VisualDSP++® 4.0
C/C++ Compiler and Library Manual
for Blackfin Processors

Revision 3.0, January 2005

Part Number
82-000410-03

Analog Devices, Inc.
One Technology Way
Norwood, Mass. 02062-9106

ANALOG DEVICES

Copyright Information

©2005 Analog Devices, Inc., ALL RIGHTS RESERVED. This document may not be reproduced in any form without prior, express written consent from Analog Devices, Inc.

Printed in the USA.

Disclaimer

Analog Devices, Inc. reserves the right to change this product without prior notice. Information furnished by Analog Devices is believed to be accurate and reliable. However, no responsibility is assumed by Analog Devices for its use; nor for any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under the patent rights of Analog Devices, Inc.

Trademark and Service Mark Notice

The Analog Devices logo, the CROSSCORE logo, VisualDSP++, Blackfin, and EZ-KIT Lite are registered trademarks of Analog Devices, Inc.

All other brand and product names are trademarks or service marks of their respective owners.

Contents

PREFACE

Purpose of This Manual xxxv
Intended Audience xxxv
Manual Contents Description xxxvi
What's New in This Manual xxxvi
Technical or Customer Support xxxvii
Supported Processors xxxviii
Product Information xxxviii
MyAnalog.com xxxviii
Processor Product Information xxxix
Related Documents xl
Online Technical Documentation xl
Accessing Documentation From VisualDSP++ xli
Accessing Documentation From Windows xli
Accessing Documentation From the Web xlii
Printed Manuals xlii
VisualDSP++ Documentation Set xlii
Hardware Tools Manuals xliii
Processor Manuals xliii

Data Sheets xliii
Notation Conventions xliv

COMPILER

C/C++ Compiler Overview 1-2
Compiler Command-Line Interface 1-4
Running the Compiler 1-5
Specifying Compiler Options in VisualDSP++ 1-9
C/C++ Compiler Command-Line Switches 1-13
C/C++ Compiler Switch Summaries 1-13
C/C++ Mode Selection Switch Descriptions 1-24
-c89 1-24
-c++ 1-24
C/C++ Compiler Common Switch Descriptions 1-25
sourcefile 1-25
-@ filename 1-25
-A name(tokens) 1-25
-alttok 1-26
-bss 1-27
-build-lib 1-27
-C 1-27
-c 1-27
-const-read-write 1-27
-cplbs 1-28
-Dmacro[=definition] 1-28

-debug-types <file.h> 1-28
-decls-<weak|strong> 1-29
-default-linkage-{asm|C|C++} 1-29
-double-size-any 1-29
-double-size-{32 | 64} 1-29
-dry 1-30
-dryrun 1-30
-E 1-30
-ED 1-30
-EE 1-30
-extra-keywords 1-31
-extra-loop-loads 1-31
-fast-fp 1-31
-flags{-asm | -compiler | -lib | -link | -mem} switch [,switch2 [,...]] 1-31
-force-circbuf 1-32
-force-link 1-32
-fp-associative 1-32
-full-io 1-33
-full-version 1-33
-g 1-33
-guard-vol-loads 1-34
-H 1-34
-HH 1-34
-h[elp] 1-34

-I directory 1-34
-I- 1-35
-i 1-36
-ieee-fp 1-36
-include filename 1-36
-ipa 1-36
-jcs2l 1-37
-jcs2l+ 1-37
-jump-<constdata|data|code> 1-37
-L directory[{,|;} directory…] 1-37
-l library 1-38
-M 1-38
-MD 1-39
-MM 1-39
-Mo filename_ 1-39
-Mt filename 1-39
-MQ 1-39
-map filename 1-40
-mem 1-40
-multicore 1-40
-multiline 1-40
-no-alttok 1-41
-no-annotate 1-41
-no-bss 1-41

-no-builtin 1-41
-no-circbuf 1-42
-no-defs 1-42
-no-extra-keywords 1-42
-no-force-link 1-42
-no-fp-associative 1-43
-no-full-io 1-43
-no-int-to-fract 1-43
-no-jcs2l 1-44
-no-jcs2l+ 1-44
-no-mem 1-44
-no-multiline 1-44
-no-saturation 1-44
-no-std-ass 1-44
-no-std-def 1-45
-no-std-inc 1-45
-no-std-lib 1-45
-no-threads 1-45
-O[0|1] 1-45
-Oa 1-46
-Ofp 1-46
-Og 1-46
-Os 1-47
-Ov num 1-47

-o filename 1-47
-P 1-47
-PP 1-47
-p[1|2] 1-47
-path {-asm | -compiler | -lib | -link | -mem} filename 1-48
-path-install directory 1-48
-path-output directory 1-48
-path-temp directory 1-48
-pch 1-49
-pchdir directory 1-49
-pedantic 1-49
-pedantic-errors 1-49
-pguide 1-50
-pplist filename 1-50
-proc processor 1-50
-progress-rep-func 1-51
-progress-rep-gen-opt 1-51
-progress-rep-mc-opt 1-52
-R directory[{;|,}directory ...] 1-52
-R- 1-52
-reserve register[, register ...] 1-53
-S 1-53
-s 1-53
-sat32 1-53

-sat40 1-53
-save-temps 1-54
-sdram 1-54
-section <id>=<section_name>[,<id>=section_name>] 1-54
-show 1-55
-signed-bitfield 1-55
-signed-char 1-55
-si-revision version 1-55
-structs-do-not-overlap 1-56
-syntax-only 1-57
-sysdefs 1-57
-T filename 1-58
-threads 1-58
-time 1-59
-U macro 1-59
-unsigned-bitfield 1-59
-unsigned-char 1-60
-v 1-60
-verbose 1-60
-version 1-60
-W{error|remark|suppress|warn} number[, number...] 1-60
-Werror-limit number 1-61
-Wremarks 1-61
-Wterse 1-61

-w 1-61
-warn-protos 1-62
-workaround <workaroundid>[,<workaroundid> …] 1-62
-write-files 1-65
-write-opts 1-66
-xref filename 1-66
C++ Mode Compiler Switch Descriptions 1-67
-anach 1-67
-eh 1-68
-no-anach 1-69
-no-demangle 1-69
-no-eh 1-69
-no-implicit-inclusion 1-69
-no-rtti 1-69
-rtti 1-70
Environment Variables Used by the Compiler 1-70
Optimization Control 1-71
Interprocedural Analysis 1-74
Interaction with Libraries 1-76
C/C++ Compiler Language Extensions 1-78
Inline Function Support Keyword (inline) 1-81
Inline Assembly Language Support Keyword (asm) 1-82
Assembly Construct Template 1-84
asm() Constructs Syntax 1-84

asm() Construct Syntax Rules 1-85
asm() Construct Template Example 1-86
Assembly Construct Operand Description 1-87
Assembly Constructs With Multiple Instructions 1-93
Assembly Construct Reordering and Optimization 1-94
Assembly Constructs With Input and Output Operands 1-94
Assembly Constructs and Flow Control 1-95
Guildelines on the Use of asm() Statements 1-96
Bank Type Qualifiers 1-96
Placement Support Keyword (section) 1-97
Placement of Compiler-Generated Data 1-98
Boolean Type Support Keywords (bool, true, false) 1-98
Pointer Class Support Keyword (restrict) 1-99
Non-Constant Aggregate Initializer Support 1-100
Indexed Initializer Support 1-100
Variable-Length Arrays 1-102
C++ Style Comments 1-102
Compiler Built-In Functions 1-103
Fractional Value Built-in Functions in C 1-104
fract16 Built-in Functions 1-106
fract32 Built-in Functions 1-108
fract2x16 Built-in Functions 1-110
ETSI Built-in Functions 1-115
ETSI Support 1-116

32-Bit Fractional ETSI Routines 1-119
16-Bit Fractional ETSI Routines 1-123
Fractional Value Built-In Functions in C++ 1-128
Fractional Literal Values in C 1-129
Complex Fractional Built-In Functions in C 1-130
Complex Operations in C++ 1-131
Viterbi History and Decoding Functions 1-132
Circular Buffer Built-In Functions 1-134
Automatic Circular Buffer Generation 1-134
Explicit Circular Buffer Generation 1-135
Circular Buffer Increment of an Index 1-135
Circular Buffer Increment of a Pointer 1-135
Endian-Swapping Intrinsics 1-136
System Built-In Functions 1-137
Compiler Performance Built-in Functions 1-139
Video Operation Built-In Functions 1-140
Function Prototypes 1-141
Example of Use: Sum of Absolute Difference 1-145
Misaligned Data Built-In Functions 1-147
Pragmas 1-148
Data Alignment Pragmas 1-149
#pragma align num 1-150
#pragma alignment_region (alignopt) 1-152
#pragma pack (alignopt) 1-153

#pragma pad (alignopt) 1-155
Interrupt Handler Pragmas 1-156
Loop Optimization Pragmas 1-156
#pragma all_aligned 1-157
#pragma different_banks 1-157
#pragma loop_count(min, max, modulo) 1-157
#pragma no_alias 1-158
#pragma no_vectorization 1-159
#pragma vector_for 1-159
General Optimization Pragmas 1-159
Linking Control Pragmas 1-160
#pragma linkage_name identifier 1-161
#pragma retain_name 1-161
#pragma section/#pragma default_section 1-162
#pragma symbolic_ref 1-164
#pragma weak_entry 1-166
Function Side-Effect Pragmas 1-167
#pragma alloc 1-167
#pragma const 1-168
#pragma noreturn 1-168
#pragma pure 1-168
#pragma regs_clobbered string 1-169
#pragma result_alignment (n) 1-173
Template Instantiation Pragmas 1-173

#pragma instantiate instance 1-174
#pragma do_not_instantiate instance 1-175
#pragma can_instantiate instance 1-175
Header File Control Pragmas 1-175
#pragma hdrstop 1-175
#pragma no_implicit_inclusion 1-176
#pragma no_pch 1-177
#pragma once 1-177
#pragma system_header 1-178
Diagnostic Control Pragmas 1-178
Modifying the Severity of Specific Diagnostics 1-178
Modifying the Behavior of an Entire Class of Diagnostics 1-179
Saving or Restoring the Current Behavior of All Diagnostics .. 1-180
Memory Bank Pragmas 1-181
#pragma code_bank(bankname) 1-182
#pragma data_bank(bankname) 1-182
#pragma stack_bank(bankname) 1-183
#pragma bank_memory_kind(bankname, kind) 1-184
#pragma bank_read_cycles(bankname, cycles) 1-185
#pragma bank_write_cycles(bankname, cycles) 1-186
#pragma bank_optimal_width(bankname, width) 1-186
GCC Compatibility Extensions 1-187
Statement Expressions 1-187
Type Reference Support Keyword (Typeof) 1-189

GCC Generalized Lvalues 1-190
Conditional Expressions With Missing Operands 1-190
Hexadecimal Floating-Point Numbers 1-191
Zero-Length Arrays 1-191
Variable Argument Macros 1-192
Line Breaks in String Literals 1-192
Arithmetic on Pointers to Void and Pointers to Functions 1-192
Cast to Union 1-193
Ranges in Case Labels 1-193
Declarations Mixed With Code 1-193
Escape Character Constant 1-194
Alignment Inquiry Keyword (__alignof__) 1-194
(asm) Keyword for Specifying Names in Generated Assembler 1-194
Function, Variable and Type Attribute Keyword (__attribute__) 1-195
Preprocessor-Generated Warnings 1-195
Blackfin Processor-Specific Functionality 1-196
Startup Code Overview 1-196
Support for argv/argc 1-197
Profiling With Instrumented Code 1-197
Generating Instrumented Code 1-198
Running the Executable 1-198
Post-Processing mon.out File 1-200
Computing Cycle Counts 1-200
Controlling Available Memory Size 1-201

Interrupt Handler Support 1-201
Defining an ISR 1-202
Registering an ISR 1-204
ISRs and ANSI C Signals 1-204
Saved Processor Context 1-205
Fetching Event Details 1-206
Fetching Saved Registers 1-207
Caching and Memory Protection 1-208
___cplb_ctrl Control Variable 1-209
CPLB Installation 1-211
Cache Configurations 1-212
Default Cache Configuration 1-213
Changing Cache Configuration 1-214
Cache Invalidation 1-214
LDFs and Cache 1-216
CPLB Replacement and Cache Modes 1-218
Cache Flushing 1-220
Using the _cplb_mgr Routine 1-221
Caching and Asynchronous Change 1-223
Migrating LDFs From Previous VisualDSP++ Installations 1-223
C/C++ Preprocessor Features 1-225
Predefined Macros 1-225
Writing Preprocessor Macros 1-228
C/C++ Run-Time Model and Environment 1-230

C/C++ Run-Time Header and Startup Code 1-232
CRT Header Overview 1-232
CRT Description 1-234
Declarations 1-235
Start and Register Settings 1-235
Event Vector Table 1-236
Stack Pointer and Frame Pointer 1-237
Cycle Counter 1-238
DAG Port Selection 1-238
Clock Speed 1-238
Memory Initialization 1-239
Device Initialization 1-240
CPLB Initialization 1-240
Enable Interrupts 1-240
Lower Processor Priority 1-241
Mark Registers 1-241
Terminate Stack Frame Chain 1-242
Profiler Initialization 1-242
C++ Constructor Invocation 1-242
Argument Parsing 1-243
Calling _main and _exit 1-243
Using Memory Sections 1-243
Using Multiple Heaps 1-245
Defining a Heap 1-245

Defining Heaps at LinkTime 1-246
Defining Heaps at RunTime 1-247
Tips for Working With Heaps 1-247
Standard Heap Interface 1-248
Using the Alternate Heap Interface 1-248
C++ Run-time Support for the Alternate Heap Interface 1-250
Freeing Space 1-250
Dedicated Registers 1-251
Call Preserved Registers 1-252
Scratch Registers 1-252
Stack Registers 1-253
Managing the Stack 1-253
Transferring Function Arguments and Return Value 1-257
Passing Arguments 1-257
Return Values 1-259
Using Data Storage Formats 1-261
Floating-Point Data Size 1-262
Floating-Point Binary Formats 1-265
IEEE Floating-Point Format 1-265
Variants of IEEE Floating-Point Support 1-267
Fract16 and Fract32 Data Representation 1-267
C/C++ and Assembly Interface 1-269
Calling Assembly Subroutines From C/C++ Programs 1-269
Calling C/C++ Functions From Assembly Programs 1-272

Using Mixed C/C++ and Assembly Naming Conventions . 1-273
Compiler C++ Template Support 1-275
Template Instantiation 1-275
Identifying Uninstantiated Templates 1-276

ACHIEVING OPTIMAL PERFORMANCE FROM C/C++ SOURCE CODE

General Guidelines 2-3
How the Compiler Can Help 2-4
Using the Compiler Optimizer 2-4
Using the Statistical Profiler 2-5
Using Profile-Guided Optimization 2-6
Using Interprocedural Optimization 2-7
Data Types 2-8
Optimizing a Struct 2-9
Bit Fields 2-11
Avoiding Emulated Arithmetic 2-12
Getting the Most From IPA 2-13
Initialize Constants Statically 2-13
Word-Aligning Your Data 2-14
Using __builtin_aligned 2-15
Avoiding Aliases 2-16
Indexed Arrays Versus Pointers 2-18
Trying Pointer and Indexed Styles 2-19
Function Inlining 2-19

Using Inline asm Statements 2-20
Memory Usage 2-21
Using the Bank Qualifier 2-23
Improving Conditional Code 2-24
Loop Guidelines 2-25
Keeping Loops Short 2-25
Avoiding Unrolling Loops 2-25
Avoiding Loop-Carried Dependencies 2-26
Avoiding Loop Rotation by Hand 2-27
Avoiding Array Writes in Loops 2-28
Inner Loops Versus Outer Loops 2-28
Avoiding Conditional Code in Loops 2-29
Avoiding Placing Function Calls in Loops 2-30
Avoiding Non-Unit Strides 2-30
Loop Control 2-31
Using the Restrict Qualifier 2-32
Using the Const Qualifier 2-33
Avoiding Long Latencies 2-34
Using Built-In Functions in Code Optimization 2-35
Fractional Data 2-35
System Support Built-In Functions 2-37
Using Circular Buffers 2-38
Smaller Applications: Optimizing for Code Size 2-40
Using Pragmas for Optimization 2-42

Function Pragmas 2-42
#pragma alloc 2-42
#pragma const 2-43
#pragma pure 2-43
#pragma result_alignment 2-44
#pragma regs_clobbered 2-44
#pragma optimize_{off|for_speed|for_space|as_cmd_line} .. 2-46
Loop Optimization Pragmas 2-47
#pragma loop_count 2-47
#pragma no_vectorization 2-47
#pragma vector_for 2-48
#pragma all_aligned 2-49
#pragma different_banks 2-50
#pragma no_alias 2-50
Useful Optimization Switches 2-51
Example: How the Optimizer Works 2-52
Assembly Optimizer Annotations 2-54
Procedure Statistics 2-55
Loop Identification 2-58
Loop Identification Annotations 2-58
Resource Definitions 2-60
File Position 2-63
Vectorization 2-65
Unroll and Jam 2-66

Loop Flattening 2-69
Vectorization Annotations 2-71

C/C++ RUN-TIME LIBRARY

C and C++ Run-Time Library Guide 3-3
Calling Library Functions 3-3
Using the Compiler's Built-In Functions 3-5
Linking Library Functions 3-5
Library Function Re-Entrancy and Multi-Threaded Environments 3-9
Support Functions for Private Data 3-11
Support Functions for Locking 3-12
Other Support Functions for Multi-Core Applications 3-12
Library Placement 3-12
Section Placement 3-13
Working With Library Header Files 3-14
assert.h 3-16
ctype.h 3-16
device.h 3-16
device_int.h 3-16
errno.h 3-16
float.h 3-17
iso646.h 3-17
limits.h 3-18
locale.h 3-18
math.h 3-18

setjmp.h 3-20
signal.h 3-20
stdarg.h 3-20
stddef.h 3-20
stdio.h 3-21
stdlib.h 3-24
string.h 3-24
time.h 3-24
Calling Library Functions from an ISR 3-26
Abridged C++ Library Support 3-27
Embedded C++ Library Header Files 3-27
complex 3-27
exception 3-28
fract 3-28
fstream 3-28
iomanip 3-28
ios 3-28
iosfwd 3-28
iostream 3-29
istream 3-29
new 3-29
ostream 3-29
shortfract 3-29
sstream 3-29

stdexcept 3-29
streambuf 3-30
string 3-30
strstream 3-30
C++ Header Files for C Library Facilities 3-30
Embedded Standard Template Library Header Files 3-31
algorithm 3-31
deque 3-31
functional 3-32
hash_map 3-32
hash_set 3-32
iterator 3-32
list 3-32
map 3-32
memory 3-32
numeric 3-32
queue 3-32
set 3-33
stack 3-33
utility 3-33
vector 3-33
fstream.h 3-33
iomanip.h 3-33
iostream.h 3-33

new.h 3-33
Using the Thread-Safe C/C++ Run-Time Libraries with VDK 3-34
File I/O Support 3-34
Extending I/O Support To New Devices 3-35
DevEntry Structure 3-35
Registering New Devices 3-40
Pre-Registering Devices 3-41
Default Device 3-42
Remove and Rename Functions 3-43
Default Device Driver Interface 3-43
Data Packing For Primitive I/O 3-44
Data Structure for Primitive I/O 3-45
Documented Library Functions 3-49
C Run-Time Library Reference 3-53
abort 3-54
abs 3-55
acos 3-56
adi_acquire_lock 3-57
adi_core_id 3-59
adi_obtain_mc_slot 3-60
asctime 3-63
asin 3-65
atan 3-66
atan2 3-67

atexit 3-68
atof 3-69
atoi 3-72
atol 3-73
atold 3-74
atoll 3-77
bsearch 3-78
cache_invalidate 3-80
calloc 3-83
ceil 3-84
clock 3-85
cos 3-86
cosh 3-88
cplb_hdr 3-89
cplb_init 3-91
cplb_mgr 3-94
ctime 3-98
difftime 3-99
disable_data_cache 3-100
div 3-101
enable_data_cache 3-102
exit 3-104
exp 3-105
fabs 3-106

floor 3-107
flush_data_cache 3-108
fmod 3-110
free 3-111
frexp 3-112
gmtime 3-113
heap_calloc 3-114
heap_free 3-116
heap_init 3-118
heap_install 3-119
heap_lookup 3-121
heap_malloc 3-123
heap_realloc 3-125
heap_space_unused 3-127
interrupt 3-128
isalnum 3-130
isalpha 3-131
iscntrl 3-132
isdigit 3-133
isgraph 3-134
isinf 3-135
islower 3-137
isnan 3-138
isprint 3-140

ispunct 3-141
isspace 3-142
isupper 3-143
isxdigit 3-144
l1_memcpy, memcpy_l1 3-145
labs 3-147
ldexp 3-148
ldiv 3-149
localtime 3-151
log 3-153
log10 3-154
longjmp 3-155
malloc 3-157
memchr 3-158
memcmp 3-159
memcpy 3-160
memmove 3-161
memset 3-162
mktime 3-163
modf 3-165
pow 3-166
qsort 3-167
raise 3-169
rand 3-171

realloc 3-172
register_handler 3-174
setjmp 3-176
signal 3-177
sin 3-179
sinh 3-182
space_unused 3-183
sqrt 3-184
srand 3-185
strcat 3-186
strchr 3-187
strcmp 3-188
strcoll 3-189
strcpy 3-190
strcspn 3-191
strerror 3-192
strftime 3-193
strlen 3-197
strncat 3-198
strncmp 3-199
strncpy 3-200
strpbrk 3-201
strrchr 3-202
strspn 3-203

strstr 3-204
strtod 3-205
strtof 3-207
strtold 3-209
strtok 3-211
strtol 3-213
strtoll 3-215
strtoul 3-217
strtoull 3-219
strxfrm 3-221
tan 3-222
tanh 3-223
time 3-224
tolower 3-225
toupper 3-226
va_arg 3-227
va_end 3-230
va_start 3-231

DSP RUN-TIME LIBRARY

DSP Run-Time Library Guide 4-2
Linking DSP Library Functions 4-2
Working With Library Source Code 4-3
DSP Header Files 4-4
complex.h – Basic Complex Arithmetic Functions 4-5

cycle_count.h – Basic Cycle Counting 4-7
cycles.h – Cycle Counting with Statistics 4-8
filter.h – Filters and Transformations 4-8
math.h – Math Functions 4-12
matrix.h – Matrix Functions 4-16
stats.h – Statistical Functions 4-22
vector.h – Vector Functions 4-24
window.h – Window Generators 4-31
Measuring Cycle Counts 4-33
Basic Cycle Counting Facility 4-33
Cycle Counting Facility with Statistics 4-35
Using time.h to Measure Cycle Counts 4-38
Determining the Processor Clock Rate 4-40
Considerations when Measuring Cycle Counts 4-41
DSP Run-Time Library Reference 4-43
a_compress 4-44
a_expand 4-45
alog 4-46
alog10 4-48
arg 4-50
autocoh 4-51
autocorr 4-53
cabs 4-55
cadd 4-56

cartesian 4-57
cdiv 4-59
cexp 4-60
cfft 4-61
cfftf 4-63
cfftrad4 4-66
cfft2d 4-68
cfir 4-70
clip 4-72
cmlt 4-73
coeff_iirdf1 4-74
conj 4-76
convolve 4-77
conv2d 4-79
conv2d3x3 4-81
copysign 4-82
cot 4-83
countones 4-84
crosscoh 4-85
crosscorr 4-87
csub 4-89
fir 4-90
fir_decima 4-92
fir_interp 4-94

gen_bartlett 4-98
gen_blackman 4-100
gen_gaussian 4-101
gen_hamming 4-102
gen_hanning 4-103
gen_harris 4-104
gen_kaiser 4-105
gen_rectangular 4-106
gen_triangle 4-107
gen_vonhann 4-109
histogram 4-110
ifft 4-112
ifftrad4 4-114
ifft2d 4-116
iir 4-118
iirdf1 4-121
max 4-125
mean 4-126
min 4-127
mu_compress 4-128
mu_expand 4-129
norm 4-130
polar 4-131
rfft 4-134

rfftrad4 4-136
rfft2d 4-138
rms 4-140
rsqrt 4-141
twidfftrad2 4-142
twidfftrad4 4-144
twidfftf_fr16 4-146
twidfft2d 4-148
var 4-150
zero_cross 4-151

ADSP-BF561 PROGRAMMING GUIDE

Dual-Core Blackfin Architecture Overview A-3
Approaches Supported in VisualDSP++ A-4
Single-Core Application A-5
Shared Memory A-6
Synchronization A-7
Cache, Startup and Events A-7
Creating Customized LDFs A-7
One Application Per Core A-8
Shared Memory A-9
Sharing Data A-9
Sharing Code A-11
Shared Code With Private Data A-11
Synchronization A-11

Cache, Startup and Events A-12
Creating Customized LDFs A-13
Single Application/Dual Core A-13
Project Conventions A-14
Multi-Core Linking A-14
Creating the LDF A-15
Shared Memory A-16
Shared Data A-16
Sharing Code in L2 Memory A-16
Synchronization A-17
Cache, Startup and Events A-17
Run-Time Library Functions A-18
Reentrancy A-18
Placement A-18
Restrictions On Dual-Core Applications A-20
Compiler Facilities A-20
Cross-Core Memory References A-20
Debugger Breakpoints A-20
Dual-Core Programming Examples A-21
Single-Core Application Example A-21
One Application Per Core Example A-23
Single Application/Dual Core Example A-27
Shared Memory, Breakpoints and printf A-29
Synchronization Functions A-31

INDEX

PREFACE

Thank you for purchasing Analog Devices development software for Blackfin® embedded media processors.

Purpose of This Manual

The *VisualDSP++ 4.0 C/C++ Compiler and Library Manual for Blackfin Processors* contains information about the C/C++ compiler and run-time libraries for Blackfin embedded processors that support a Media Instruction Set Computing (MISC) architecture. This architecture is the natural merging of RISC, media functions, and signal processing characteristics that delivers signal processing performance in a microprocessor-like environment.

Intended Audience

The primary audience for this manual is a programmer who is familiar with Analog Devices' Blackfin processors. This manual assumes that the audience has a working knowledge of the Blackfin processors' architecture and instruction set and the C/C++ programming languages.

Programmers who are unfamiliar with Blackfin processors can use this manual, but they should supplement it with other texts (such as the appropriate hardware reference and instruction set reference) that provide information about your Blackfin processor architecture and instructions).

Manual Contents Description

This manual contains:

- Chapter 1, "Compiler"
  Provides information on compiler options, language extensions, C/C++/assembly interfacing, and support for C++ templates
- Chapter 2, "Achieving Optimal Performance from C/C++ Source Code"
  Shows how to optimize compiler operation.
- Chapter 3, "C/C++ Run-Time Library"
  Shows how to use library functions and provides a complete C/C++ library function reference
- Chapter 4, "DSP Run-Time Library"
  Shows how to use DSP library functions and provides a complete DSP library function reference
- Appendix A, "ADSP-BF561 Programming Guide"
  Provides various approaches and programming guidance for developing systems on ADSP-BF561 Blackfin processors

What's New in This Manual

This edition of the *VisualDSP++ 4.0 C/C++ Compiler and Library Manual for Blackfin Processors* documents support for all Blackfin processors.

Refer to *VisualDSP++ 4.0 Product Bulletin* for information on all new and updated features and other release information.

Technical or Customer Support

You can reach Analog Devices, Inc. Customer Support in the following ways:

- Visit the Embedded Processing and DSP products Web site at
  `http://www.analog.com/processors/technicalSupport`
- E-mail tools questions to
  `dsptools.support@analog.com`
- E-mail processor questions to
  `dsp.support@analog.com`
- Phone questions to **1-800-ANALOGD**
- Contact your Analog Devices, Inc. local sales office or authorized distributor
- Send questions by mail to:

```
Analog Devices, Inc.
One Technology Way
P.O. Box 9106
Norwood, MA 02062-9106
USA
```

Supported Processors

The name "*Blackfin*" refers to a family of 16-bit, embedded processors. VisualDSP++ currently supports the following Blackfin processors:

ADSP-BF531	ADSP-BF532 (formerly ADSP-21532)
ADSP-BF533	ADSP-BF535 (formerly ADSP-21535)
ADSP-BF534	ADSP-BF536
ADSP-BF537	ADSP-BF538
ADSP-BF539	ADSP-BF561
ADSP-BF566	AD6532

Product Information

You can obtain product information from the Analog Devices Web site, from the product CD-ROM, or from the printed publications (manuals).

Analog Devices is online at `www.analog.com`. Our Web site provides information about a broad range of products—analog integrated circuits, amplifiers, converters, and digital signal processors.

MyAnalog.com

`MyAnalog.com` is a free feature of the Analog Devices Web site that allows customization of a Web page to display only the latest information on products you are interested in. You can also choose to receive weekly e-mail notifications containing updates to the Web pages that meet your interests. `MyAnalog.com` provides access to books, application notes, data sheets, code examples, and more.

Registration

Visit www.myanalog.com to sign up. Click **Register** to use MyAnalog.com. Registration takes about five minutes and serves as a means to select the information you want to receive.

If you are already a registered user, just log on. Your user name is your e-mail address.

Processor Product Information

For information on embedded processors and DSPs, visit our Web site at www.analog.com/processors, which provides access to technical publications, data sheets, application notes, product overviews, and product announcements.

You may also obtain additional information about Analog Devices and its products in any of the following ways.

- E-mail questions or requests for information to
  dsp.support@analog.com
- Fax questions or requests for information to
  1-781-461-3010 (North America)
  089/76 903-557 (Europe)
- Access the FTP Web site at
  ftp ftp.analog.com or ftp 137.71.23.21
  ftp://ftp.analog.com

Related Documents

For information on product related development software, see these publications:

- *VisualDSP++ 4.0 Getting Started Guide*
- *VisualDSP++ 4.0 User's Guide*
- *VisualDSP++ 4.0 Assembler and Preprocessor Manual*
- *VisualDSP++ 4.0 Linker and Utilities Manual*
- *VisualDSP++ 4.0 Loader Manual*
- *VisualDSP++ 4.0 Product Release Bulletin*
- *VisualDSP++ Kernel (VDK) User's Guide*
- *Quick Installation Reference Card*

For hardware information, refer to your processors's hardware reference, programming reference, or data sheet. All documentation is available online. Most documentation is available in printed form.

Visit the Technical Library Web site to access all processor and tools manuals and data sheets:

`http://www.analog.com/processors/resources/technicalLibrary`

Online Technical Documentation

Online documentation includes the VisualDSP++ Help system, software tools manuals, hardware tools manuals, processor manuals, Dinkum Abridged C++ library, and Flexible License Manager (FlexLM) network license manager software documentation. You can easily search across the entire VisualDSP++ documentation set for any topic of interest using the Search function of VisualDSP++ Help system. For easy printing, supplementary `.PDF` files of most manuals are also provided.

Each documentation file type is described as follows.

File	Description
.CHM	Help system files and manuals in Help format
.HTM or .HTML	Dinkum Abridged C++ library and FlexLM network license manager software documentation. Viewing and printing the .HTML files requires a browser, such as Internet Explorer 4.0 (or higher).
.PDF	VisualDSP++ and processor manuals in Portable Documentation Format (PDF). Viewing and printing the .PDF files requires a PDF reader, such as Adobe Acrobat Reader (4.0 or higher).

Access the online documentation from the VisualDSP++ environment, Windows® Explorer, or the Analog Devices Web site.

Accessing Documentation From VisualDSP++

From the VisualDSP++ environment:

- Access VisualDSP++ online Help from the Help menu's **Contents**, **Search**, and **Index** commands.
- Open online Help from context-sensitive user interface items (toolbar buttons, menu commands, and windows).

Accessing Documentation From Windows

In addition to any shortcuts you may have constructed, there are many ways to open VisualDSP++ online Help or the supplementary documentation from Windows.

Help system files (.CHM) are located in the Help folder of VisualDSP++ environment. The .PDF files are located in the Docs folder of your VisualDSP++ installation CD-ROM. The Docs folder also contains the Dinkum Abridged C++ library and the FlexLM network license manager software documentation.

Using Windows Explorer

- Double-click the `vdsp-help.chm` file, which is the master Help system, to access all the other `.CHM` files.
- Open your VisualDSP++ installation CD-ROM and double-click any file that is part of the VisualDSP++ documentation set.

Using the Windows Start Button

- Access VisualDSP++ online Help by clicking the **Start** button and choosing **Programs**, **Analog Devices**, **VisualDSP++**, and **VisualDSP++ Documentation**.

Accessing Documentation From the Web

Download manuals in PDF format at the following Web site:
`http://www.analog.com/processors/resources/technicalLibrary/manuals`

Select a processor family and book title. Download archive (`.ZIP`) files, one for each manual. Use any archive management software, such as WinZip, to decompress downloaded files.

Printed Manuals

For general questions regarding literature ordering, call the Literature Center at **1-800-ANALOGD** (**1-800-262-5643**) and follow the prompts.

VisualDSP++ Documentation Set

To purchase VisualDSP++ manuals, call **1-603-883-2430**. The manuals may be purchased only as a kit.

If you do not have an account with Analog Devices, you are referred to Analog Devices distributors. For information on our distributors, log onto `http://www.analog.com/salesdir/continent.asp`.

Hardware Tools Manuals

To purchase EZ-KIT Lite® and In-Circuit Emulator (ICE) manuals, call **1-603-883-2430**. The manuals may be ordered by title or by product number located on the back cover of each manual.

Processor Manuals

Hardware reference and instruction set reference manuals may be ordered through the Literature Center at **1-800-ANALOGD** (**1-800-262-5643**), or downloaded from the Analog Devices Web site. Manuals may be ordered by title or by product number located on the back cover of each manual.

Data Sheets

All data sheets (preliminary and production) may be downloaded from the Analog Devices Web site. Only production (final) data sheets (Rev. 0, A, B, C, and so on) can be obtained from the Literature Center at **1-800-ANALOGD** (**1-800-262-5643**); they also can be downloaded from the Web site.

To have a data sheet faxed to you, call the Analog Devices Faxback System at **1-800-446-6212**. Follow the prompts and a list of data sheet code numbers will be faxed to you. If the data sheet you want is not listed, check for it on the Web site.

Notation Conventions

Text conventions used in this manual are identified and described as follows.

Example	Description
Close command (**File** menu)	Titles in reference sections indicate the location of an item within the VisualDSP++ environment's menu system (for example, the **Close** command appears on the **File** menu).
`{this | that}`	Alternative required items in syntax descriptions appear within curly brackets and separated by vertical bars; read the example as `this` or `that`. One or the other is required.
`[this | that]`	Optional items in syntax descriptions appear within brackets and separated by vertical bars; read the example as an optional `this` or `that`.
`[this,…]`	Optional item lists in syntax descriptions appear within brackets delimited by commas and terminated with an ellipse; read the example as an optional comma-separated list of `this`.
`.SECTION`	Commands, directives, keywords, and feature names are in text with `letter gothic` font.
`filename`	Non-keyword placeholders appear in text with italic style format.
	Note: For correct operation, ... A Note provides supplementary information on a related topic. In the online version of this book, the word **Note** appears instead of this symbol.
	Caution: Incorrect device operation may result if ... **Caution:** Device damage may result if ... A Caution identifies conditions or inappropriate usage of the product that could lead to undesirable results or product damage. In the online version of this book, the word **Caution** appears instead of this symbol.
	Warning: Injury to device users may result if ... A Warning identifies conditions or inappropriate usage of the product that could lead to conditions that are potentially hazardous for devices users. In the online version of this book, the word **Warning** appears instead of this symbol.

Additional conventions, which apply only to specific chapters, may appear throughout this document.

1 COMPILER

The C/C++ compiler (`ccblkfn.exe`) is part of Analog Devices development software for Blackfin processors.

The code examples in this manual have been compiled using VisualDSP++ 4.0. The examples compiled with other versions of VisualDSP++ may result in build errors or different output although the highlighted algorithms stand and should continue to stand in future releases of VisualDSP++.

This chapter contains:

- "C/C++ Compiler Overview" on page 1-2
  provides an overview of the C/C++ compiler for Blackfin processors.
- "Compiler Command-Line Interface" on page 1-4
  describes the operation of the compiler as it processes programs, including input and output files and command-line switches.
- "C/C++ Compiler Language Extensions" on page 1-78
  describes the `ccblkfn` compiler's extensions to the ANSI/ISO standard for the C and C++ languages.
- "Blackfin Processor-Specific Functionality" on page 1-196
  contains information that is specific to Blackfin processors only.
- "C/C++ Preprocessor Features" on page 1-225
  contains information on the preprocessor and ways to modify source compilation.

- "C/C++ Run-Time Model and Environment" on page 1-230 contains reference information about implementation of C/C++ programs, data, and function calls in Blackfin processors.
- "C/C++ and Assembly Interface" on page 1-269 describes how to call an assembly language subroutine from within a C or C++ program, and how to call a C or C++ function from within an assembly language program.
- "Compiler C++ Template Support" on page 1-275 describes how templates are instantiated at compile time

C/C++ Compiler Overview

The C/C++ compiler is designed to aid your DSP project development efforts by:

- Processing C and C++ source files, producing machine-level versions of the source code and object files
- Providing relocatable code and debugging information within the object files
- Providing relocatable data and program memory segments for placement by the linker in the processors' memory

Using C/C++, developers can significantly decrease time-to-market since it gives them the ability to efficiently work with complex signal processing data types. It also allows them to take advantage of specialized signal processing operations without having to understand the underlying processor architecture.

The C/C++ compiler compiles ANSI/ISO standard C and C++ code to support signal data processing. Additionally, Analog Devices includes within the compiler a number of C language extensions designed to assist in DSP development. The `ccblkfn` compiler runs from the VisualDSP++ environment or from an operating system command line.

The C/C++ compiler processes your C and C++ language source files and produces Blackfin assembler source files. The assembler source files are assembled by the Blackfin processor assembler (`easmblkfn`). The assembler creates Executable and Linkable Format (ELF) object files that can be linked (using the linker) to create a Blackfin processor executable file or included in an archive library using the librarian tool (`elfar`). The way in which the compiler controls the assemble, link, and archive phases of the process depends on the source input files and the compiler options used.

Your source files contain the C/C++ program to be processed by the compiler. The `ccblkfn` compiler supports the ANSI/ISO standard definitions of the C and C++ languages. For information on the C language standard, see any of the many reference texts on the C language. Analog Devices recommends the Bjarne Stroustrup text "*The C++ Programming Language*" from Addison Wesley Longman Publishing Co (ISBN: 0201889544) (1997) as a reference text for the C++ programming language.

The `ccblkfn` compiler supports a set of C/C++ language extensions. These extensions support hardware features of the Blackfin processors. For information on these extensions, see "C/C++ Compiler Language Extensions" on page 1-78.

You can set the compiler options from the **Compile** page of the **Project Options** dialog box of the VisualDSP++ Integrated Development and Debug Environment (IDDE) (see "Specifying Compiler Options in VisualDSP++" on page 1-9). These selections control how the compiler processes your source files, letting you select features that include the language dialect, error reporting, and debugger output.

For more information on the VisualDSP++ environment, see the *VisualDSP++ 4.0 User's Guide* and online Help.

Compiler Command-Line Interface

This section describes how the `ccblkfn` compiler is invoked from the command line, the various types of files used by and generated from the compiler, and the switches used to tailor the compiler's operation.

This section contains:

- "Running the Compiler" on page 1-5
- "Specifying Compiler Options in VisualDSP++" on page 1-9
- "C/C++ Compiler Command-Line Switches" on page 1-13
- "Environment Variables Used by the Compiler" on page 1-70
- "Optimization Control" on page 1-71

By default, the compiler runs with Analog Extensions for C code enabled. This means that the compiler processes source files written in ANSI/ISO standard C language supplemented with Analog Devices extensions. Table 1-1 on page 1-6 lists valid extensions of source files the compiler operates upon. By default, the compiler processes the input file through the listed stages to produce a `.DXE` file (see file names in Table 1-2 on page 1-8). Table 1-3 on page 1-13 lists the switches that select the language dialect.

Although many switches are generic between C and C++, some of them are valid in C++ mode only. A summary of the generic C/C++ compiler switches appears in Table 1-4 on page 1-14. A summary of the C++-specific compiler switches appears in Table 1-5 on page 1-23. The summaries are followed by descriptions of each switch.

When developing a DSP project, sometimes it is useful to modify the compiler's default options settings. The way the compiler's options are set depends on the environment used to run the DSP development software. See "Specifying Compiler Options in VisualDSP++" on page 1-9 for more information.

Running the Compiler

Use the following syntax for the `ccblkfn` command line

```
ccblkfn [-switch [-switch …] sourcefile [sourcefile …]]
```

to run the compiler with

`ccblkfn`	Name of the compiler program for Blackfin processors.
`-switch`	Switch (or switches) to process. The compiler has many switches. These switches select the operations and modes for the compiler and other tools. Command-line switches are case-sensitive. For example, `-O` is not the same as `-o`.
`sourcefile`	Name of the file to be preprocessed, compiled, assembled, and/or linked

A file name can include the drive, directory, file name, and file extension. The compiler supports both Win32- and POSIX-style paths, using either forward or back slashes as the directory delimiter. It also supports UNC path names (starting with two slashes and a network name).

The `ccblkfn` compiler uses the file extension to determine what the file contains and what operations to perform upon it. Table 1-2 on page 1-8 lists the allowed extensions.

For example, the following command line

```
ccblkfn -proc ADSP-BF535 -O -Wremarks -o program.dxe source.c
```

runs `ccblkfn` with

`-proc ADSP-BF535`	Specifies compiler instructions unique to the ADSP-BF535 processor
`-O`	Specifies optimization for the compiler
`-Wremarks`	Selects extra diagnostic remarks in addition to warning and error messages

`-o program.dxe`	Selects a name for the compiled, linked output
`source.c`	Specifies the C language source file to be compiled

The following example command line for the C++ mode,

```
ccblkfn -proc ADSP-BF535 -c++ source.cpp
```

runs `ccblkfn` with

`-c++`	Specifies all of the source files to be compiled in C++ mode
`source.cpp`	Specifies the C++ language source file to be compiled

The normal function of `ccblkfn` is to invoke the compiler, assembler, and linker as required to produce an executable object file. The precise operation is determined by the extensions of the input file names and by various switches.

In normal operation, the compiler uses the files listed in Table 1-1 to perform a specified action.

Table 1-1. File Extensions Specifying Compiler Action

Extension	Action
`.c .cpp .cxx .cc .c++`	Source file is compiled, assembled, and linked.
`.asm`, `.dsp`, or `.s`	Assembly language source file is assembled and linked.
`.doj`	Object file (from previous assembly) is linked.

If multiple files are specified, each is processed to produce an object file and then all the object files are presented to the linker.

You can stop this sequence at various points using appropriate compiler switches, or selecting options with the VisualDSP++ IDDE. These switches are `-E`, `-P`, `-M`, `-H`, `-S`, `-c`.

Many of the compiler's switches take a file name as an optional parameter. If you do not use the optional output name switch, `ccblkfn` names the output for you. Table 1-2 lists the type of files, names, and extensions `ccblkfn` appends to output files.

File extensions vary by command-line switch and file type. These extensions are influenced by the program that is processing the file. They search directories that you select and path information that you include in the file name. Table 1-2 indicates the extensions that the preprocessor, compiler, assembler, and linker support. The compiler supports relative and absolute directory names to define file extension paths. For information on additional search directories, see the command-line switch that controls the specific type of extensions.

When providing an input or output file name as an optional parameter, use the following guidelines.

- Use a file name (include the file extension) with either an unambiguous relative path or an absolute path. A file name with an absolute path includes the drive, directory, file name, and file extension.

  Enclose long file names within straight quotes; for example, "`long file name.c`". The compiler uses the file extension convention listed in Table 1-2 to determine the input file type.

- Verify the compiler is using the correct file. If you do not provide the complete file path as part of the parameter or add additional search directories, `ccblkfn` looks for input in the current directory.

Using the verbose output switches for the preprocessor, compiler, assembler, and linker causes each of these tools to echo the name of each file as it is processed.

The compiler refers to a number of environment variables during its operation, and these environment variables can affect the compiler's behavior. Refer to "Environment Variables Used by the Compiler" on page 1-70 for more information.

Table 1-2. Input and Output File Extensions

File Extension	File Extension Description
.c	C source file
.cpp .cxx .cc .c++	C++ source file
.h	Header file (referenced by an #include statement)
.hpp .hh .hxx .h++	C++ header file (referenced by a #include statement)
.ii, .ti	Template instantiation files – used internally by the compiler when instantiating templates
.ipa, .opa	Interprocedural analysis files – used internally by the compiler when performing interprocedural analysis.
.pgo	Execution profile generated by Profile-Guided Optimization.
.i	Preprocessed source file — created when preprocess only is specified
.s, .asm	Assembly language source files
.is	Preprocessed assembly language source — retained when -save-temps is specified
.ldf	Linker Description File
.pch	Precompiled header file.
.doj	Object file to be linked
.dlb	Library of object files to be linked as needed
.dxe	Executable file produced by compiler
.xml	Processor memory map file output
.sym	Processor symbol map file output

Specifying Compiler Options in VisualDSP++

In the VisualDSP++ Integrated Development and Debug Environment (IDDE), use the **Compile** option from the **Project Options** dialog box to access compiler functional options, shown in Figure 1-1.

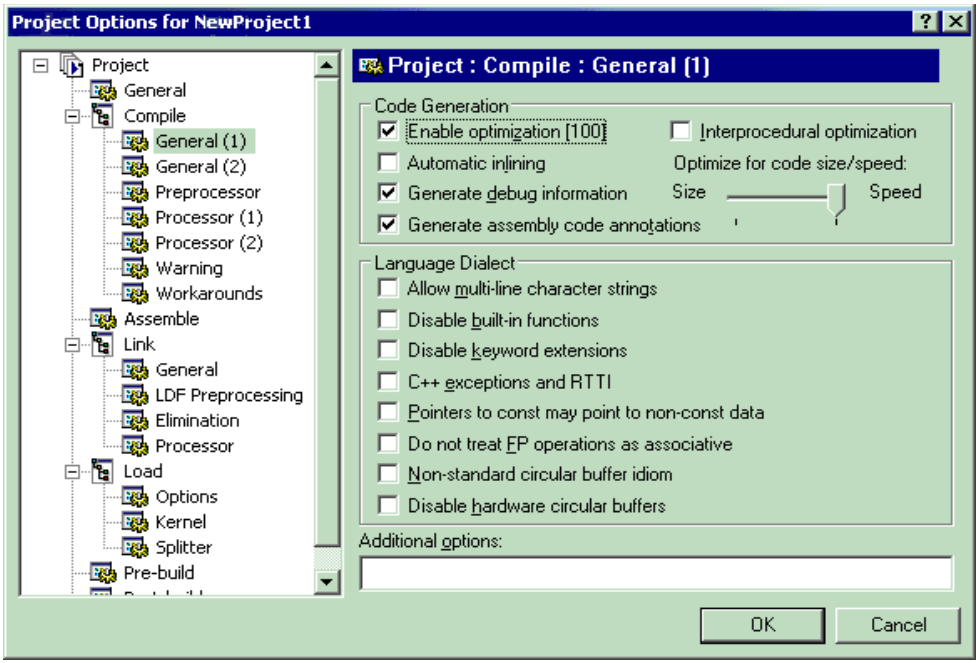

Figure 1-1. Project Options — Compile (General1) Property Page

There are several sub-pages you can access—**General1**, **General1**, **Preprocessor**, **Processor1**, **Processor2**, **Warning**, and **Workarounds**. Most page options have a corresponding compiler command-line switch described in "C/C++ Compiler Command-Line Switches" on page 1-13.

The **Additional options** field in each sub-page is used to enter the appropriate file names and options that do not have corresponding controls on the **Compile** page but are available as compiler switches.

For example, Figure 1-2 shows the first (of two) page with processor-specific information.

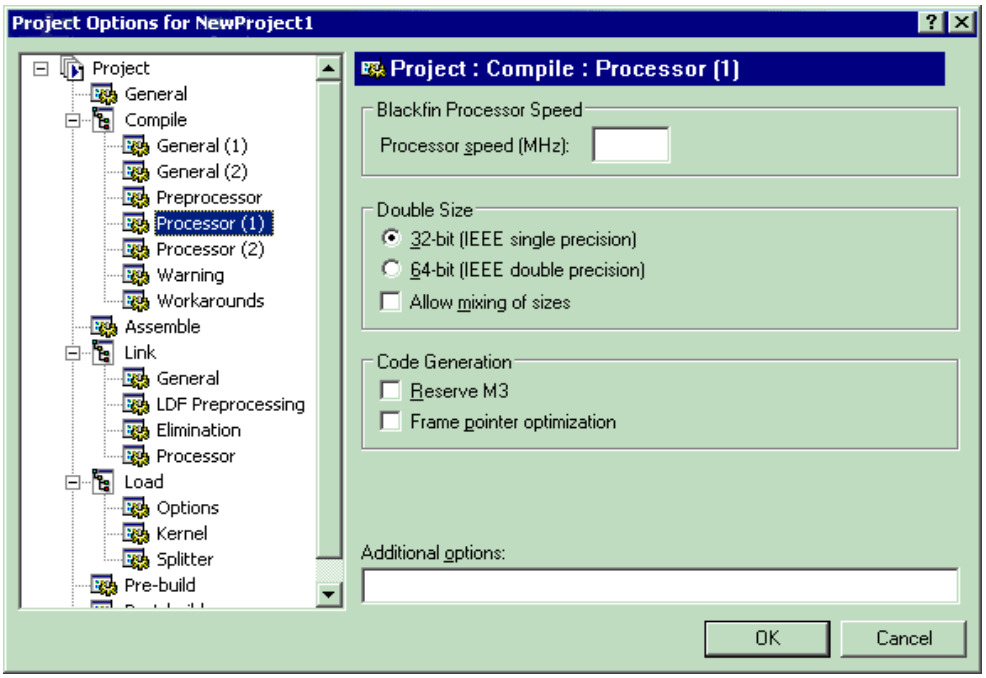

Figure 1-2. Project Options — Compile (Processor1) Property Page

Figure 1-3 shows the second (of two) page with processor-specific information.

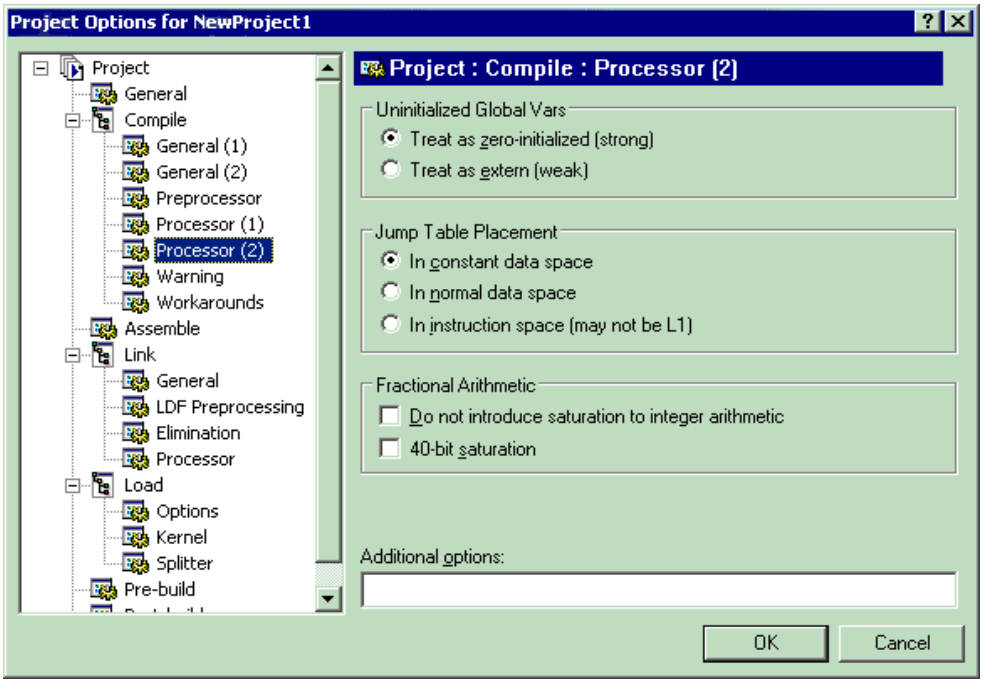

Figure 1-3. Project Options — Compile (Processor2) Property Page

Figure 1-4 shows a list of current code generator workarounds for specific hardware errata (for more information, see the switch "-workaround <workaroundid>[,<workaroundid> …]" on page 1-62.

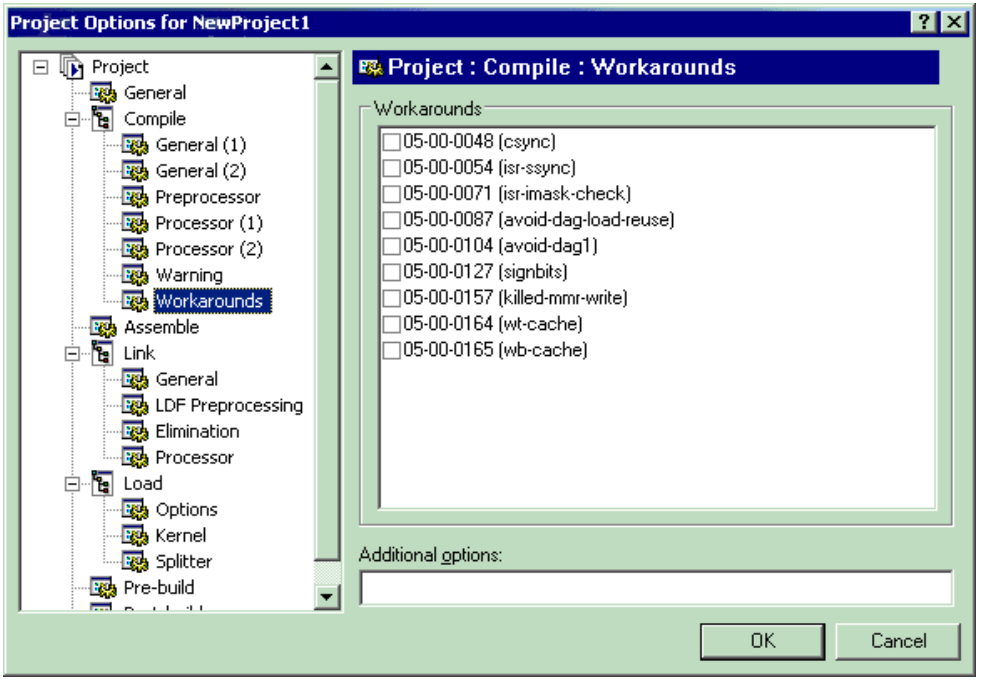

Figure 1-4. Project Options — Compile (Workarounds) Property Page

For more information, use the VisualDSP++ context-sensitive online Help for each target architecture to select information on compiler options you can specify in VisualDSP++. To do that, click on the **?** button and then click in a field or box you need information about.

C/C++ Compiler Command-Line Switches

This section describes command-line switches used when compiling. It contains a set of tables that provides a brief description of each switch. These tables are organized by type of switch. Following these tables are sections that provide detailed switch descriptions.

C/C++ Compiler Switch Summaries

This section contains a set of tables that summarize generic and specific switches (options).

- Table 1-3, "C or C++ Mode Selection Switches"
- Table 1-4, "C/C++ Compiler Common Switches"
- Table 1-5, "C++ Mode Compiler Switches" on page 1-23

A brief description of each switch appears in the sections beginning on page 1-24.

Table 1-3. C or C++ Mode Selection Switches

Switch Name	Description
`-c89` on page 1-24	Supports programs that conform to the ISO/IEC 9899:1990 standard
`-c++` on page 1-24	Supports ANSI/ISO standard C++ with Analog Devices extensions

Table 1-4. C/C++ Compiler Common Switches

Switch Name	Description
`sourcefile` on page 1-25	Specifies the file to be compiled
`-@ filename` on page 1-25	Reads command-line input from the file
`-A symbol(tokens)` on page 1-25	Asserts the specified name as a predicate
`-alttok` on page 1-26	Allows alternative keywords and sequences in sources
`-bss` on page 1-27	Causes the compiler to put global zero-initialized data into a separate BSS-style section. Set by default.
`-build-lib` on page 1-27	Directs the librarian to build a library file
`-C` on page 1-27	Retains preprocessor comments in the output file
`-c` on page 1-27	Compiles and/or assembles only, but does not link
`-const-read-write` on page 1-28	Specifies that data accessed via a pointer to `const` data may be modified elsewhere
`-cplbs` on page 1-28	Instructs the compiler to assume that CPLBs are active
`-Dmacro[=def]` on page 1-28	Defines `macro`
`-debug-types` on page 1-28	Supports building a `*.h` file directly and writing a complete set of debugging information for the header file
`-decls-<weak|strong>` on page 1-29	Determines whether uninitialised global variables should be treated as definitions or declarations
`-default-linkage-<asm|C|C++>` on page 1-29	Sets the default linkage type
`-double-size-any` (on page 1-29)	Indicates that the resulting object can be linked with objects built with any `double` size

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
`-double-size-{32|64}` (on page 1-29)	Selects 32- or 64-bit IEEE format for `double`. `-double-size-32` is the default mode.
`-dry` on page 1-30	Displays, but does not perform, main driver actions (verbose dry run)
`-dryrun` on page 1-30	Displays, but does not perform, top-level driver actions (terse dry run)
`-E` on page 1-30	Preprocesses, but does not compile, the source file
`-ED` on page 1-30	Preprocesses and sends all output to a file
`-EE` on page 1-30	Preprocesses and compiles the source file
`-extra-keywords` on page 1-31	Recognizes Blackfin processor extensions to ANSI/ISO standards for C (default mode)
`-extra-loop-loads` on page 1-31	Allows the compiler to read off the start or end of memory areas, within loops, to aid performance
`-fast-fp` on page 1-31	Links with the high-speed floating-point emulation library
`-flags` *`tool`* on page 1-31	Passes command-line switches through the compiler to other build tools
`-force-circbuf` on page 1-31	Treats array references of the form `array[i%n]` as circular buffer operations
`-force-link` on page 1-31	Forces stack frame creation for leaf functions. Always creates a new stack frame for leaf functions (defaults to ON with `-g` option set, enforced for the `-p` option)
`-fp-associative` on page 1-32	Treats floating-point multiplication and addition as associative operations
`-full-io` on page 1-33	Links with a third party, proprietary I/O library

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
`-full-version` on page 1-33	Displays the version number of the driver and processes invoked by the driver
`-g` on page 1-33	Generates DWARF-2 debug information
`-guard-vol-loads` on page 1-34	Disables interrupts during volatile loads
`-H` on page 1-34	Outputs a list of included header files, but does not compile
`-HH` on page 1-34	Outputs a list of included header files and compiles.
`-h[elp]` on page 1-34	Outputs a list of command-line switches with brief syntax descriptions
`-I` *`directory`* on page 1-34	Appends *`directory`* to the standard search path
`-I-` on page 1-34	Specifies the point in the `include` directory list where the search for header files enclosed in angle brackets should begin
`-i` on page 1-36	Outputs only header details or makefile dependencies for `include` files specified in double quotes.
`-ieee-fp` on page 1-36	Links with the fully-compliant floating-point emulation library
`-include` *`filename`* on page 1-36	Includes named file prior to each source file
`-ipa` on page 1-36	Specifies that interprocedural analysis should be performed for optimization between translation units
`-jcs2l` on page 1-37	Enables the conversion of short jumps to long jumps when necessary
`-jcs2l+` on page 1-37	Enables the conversion of `short` jumps to `long` jumps when necessary but uses the P1 register for indirect jumps when long jumps are insufficient (enabled by default)

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
`-jump-<constdata|data|code>` on page 1-37	Determines which section the compiler uses to store jump-tables that are generated for C "switch" statements.
`-L directory` on page 1-37	Appends *directory* to the standard library search path
`-l library` on page 1-38	Searches *library* for functions when linking
`-M` on page 1-38	Generates `make` rules only, but does not compile
`-MD` on page 1-39	Generates `make` rules, compiles, and prints to a file
`-MM` on page 1-39	Generates `make` rules and compiles
`-Mo filename` on page 1-39	Writes dependency information to *filename*. This switch is used in conjunction with the `-ED` or `-MD` options.
`-Mt filename` on page 1-39	Makes dependencies, where the target is renamed as *filename*
`-MQ` on page 1-39	Generates make rules only; does not compile. No notification when input files are missing.
`-map filename` on page 1-40	Directs the linker to generate a memory map of all symbols
`-mem` on page 1-40	Causes the compiler to invoke the Memory Initializer after linking the executable file
`-multicoree` on page 1-40	Selects library versions suitable for use in a multi-core environment
`-multiline` on page 1-40	Enables string literals over multiple lines (default)
`-no-alttok` on page 1-41	Does not allow alternative keywords and sequences in sources
`-no-annotate` (on page 1-41)	Disables the annotation of assembly files

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
-no-bss on page 1-41	Causes the compiler to group global zero-initialized data into the same section as global data with non-zero initializers
-no-builtin on page 1-41	Disable recognition of __builtin functions
-no-circbuf on page 1-42	Disables the automatic generation of circular buffering code
-no-defs on page 1-42	Disables preprocessor definitions: macros, include directories, library directories, run-time headers, or keyword extensions
-no-extra-keywords on page 1-42	Does not define language extension keywords which could be valid C/C++ identifiers
-no-force-link on page 1-43	Does not create a new stack frame for leaf functions, if one can be omitted. Overrides the default for -g.
-no-fp-associative (on page 1-43)	Does not treat floating-point multiplication and addition as associative operations
-no-full-io (on page 1-43)	Links with the Analog Devices I/O library. Enabled by default.
-no-int-to-fract on page 1-43	Prevents the compiler from turning integer into fractional arithmetic
-no-jcs2l on page 1-44	Disables the conversion of short jumps to long jumps
-no-jcs2l+ on page 1-44	Disables the conversion of short jumps to long jumps
-no-mem on page 1-44	Causes the compiler to not invoke the Memory Initializer after linking. Set by default.
-no-multiline on page 1-44	Disables multiple line string literal support
-no-saturation on page 1-44	Causes the compiler not to introduce saturation semantics when optimizing expressions

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
`-no-std-ass` on page 1-44	Prevents the compiler from defining standard assertions
`-no-std-def` on page 1-45	Disables normal macro definitions and also ADI keyword extensions that do not have leading underscores (__)
`-no-std-inc` on page 1-45	Searches only for preprocessor `include` header files in the current directory and in directories specified with the `-I` switch
`-no-std-lib` on page 1-45	When linking, searches for only those library files specified with the `-l` switch
`-no-threads` on page 1-45	Specifies that no support is required for multithreaded applications
`-O[0|1]` on page 1-45	Enables code optimizations
`-Oa` on page 1-46	Enables automatic function inlining
`-Ofp` on page 1-46	Offsets the Frame Pointer to allow more short load and store instructions. Reduces debugger capabilities, when used with `-g`.
`-Og` (on page 1-46)	Enables a compiler mode that performs optimizations while still preserving the debugging information
`-Os` on page 1-47	Optimizes the file to decrease code size
`-Ov num` on page 1-47	Controls speed versus size optimizations
`-o` *filename* on page 1-47	Specifies the output file name
`-P` on page 1-47	Preprocesses, but does not compile, the source file; output does not contain `#line` directives
`-PP` on page 1-47	Preprocesses and compiles the source file; output does not contain `#line` directives.

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
`-p[1|2]` on page 1-47	Generates profiling instrumentation
`-path-{asm|compiler|lib|link|mem}` *filename* on page 1-48	Uses the specified directory as the location of the specified compilation tool (assembler, compiler, library builder, linker, or memory initializer)
`-path-install` *directory* on page 1-48	Uses the specified directory as the location of all compilation tools
`-path-output` *directory* on page 1-48	Specifies the location of non-temporary files
`-path-temp` *directory* on page 1-48	Specifies the location of temporary files
`-pch` on page 1-49	Enables automatic generation and use of precompiled header files
`-pchdir` *directory* on page 1-49	Specifies an alternative directory to `PCHRepository` in which to store precompiled header files
`-pedantic` on page 1-49	Issues compiler warnings for constructs that are not strictly ISO/ANSI standard C/C++ compliant
`-pedantic-errors` on page 1-49	Issues compiler errors for constructs that are not strictly ISO/ANSI standard C/C++ compliant
`-pguide` on page 1-50	Adds instrumentation for the gathering of a profile as the first stage of performing profile-guided optimization
`-pplist` *filename* on page 1-50	Outputs a raw preprocessed listing to the specified file
`-proc` *processor* on page 1-50	Specifies a processor for which the compiler should produce suitable code
`-progress-rep-func` on page 1-51	Issues a diagnostic message each time the compiler starts compiling a new function. Equivalent to `-Wwarn=cc1472`.
`-progress-rep-gen-opt` on page 1-51	Issues a diagnostic message each time the compiler starts a new generic optimization pass on the current function. Equivalent to `-Wwarn=cc1473`.

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
`-progress-rep-mc-opt` on page 1-52	Issues a diagnostic message each time the compiler starts a new machine-specific optimization pass on the current function. Equivalent to `-Wwarn=cc1474`.
`-R` *`directory`* on page 1-52	Appends *`directory`* to the standard search path for source files
`-R-` on page 1-52	Removes all directories from the source file search directory list
`-reserve <reg1>[,reg2...]` on page 1-53	Reserves certain registers from compiler use. **Note:** Reserving registers can have a detrimental effect on the compiler's optimization capabilities.
`-S` on page 1-53	Stops compilation before running the assembler
`-s` on page 1-53	When linking, removes debugging information from the output executable file
`-sat32` on page 1-53	Saturates all accumulations at 32 bits, which is the default
`-sat40` on page 1-53	Saturates all accumulations at 40 bits rather than the default 32 bits
`-save-temps` on page 1-54	Saves intermediate files
`-sdram` on page 1-54	Instructs the compiler to assume that at least bank 0 of external SDRAM will be present and enabled
`-section <id>=<section name>` on page 1-54	Orders the compiler to place data/program of type "`id`" into the section "`section_name`"
`-show` on page 1-55	Displays the driver command-line information
`-signed-bitfield` on page 1-55	Makes the default type for `int bitfields` signed
`-signed-char` on page 1-55	Makes the default type for `char` signed

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
`-si-revision` *version* on page 1-55	Specifies a silicon revision of the specified processor. The default setting is the latest silicon revision
`-structs-do-not-overlap` on page 1-56	Specifies that `struct` copies may use "`memcpy`" semantics, rather than the usual "`memmove`" behavior
`-syntax-only` on page 1-57	Checks the source code for compiler syntax errors, but does not write any output
`-sysdefs` on page 1-57	Instructs the driver to define preprocessor macros that describe the current user and machine
`-T` *filename* on page 1-58	Specifies the Linker Description File
`-threads` on page 1-58	Enables the support for multithreaded applications
`-time` on page 1-59	Displays the elapsed time as part of the output information on each part of the compilation process
`-U`*macro* on page 1-59	Undefines *macro*
`-unsigned-bitfield` on page 1-60	Makes the default type for plain `int` bitfields unsigned
`-unsigned-char` on page 1-60	Makes the default type for `char` unsigned
`-v` on page 1-60	Displays version and command-line information for all compilation tools
`-verbose` on page 1-60	Displays command-line information for all compilation tools as they process each file
`-version` on page 1-60	Displays version information for all compilation tools as they process each file
`-Werror` *number* on page 1-61	Overrides the default severity of the specified messages (errors, remarks, or warnings)
`-Werror-limit` *number* on page 1-61	Stops compiling after reaching the specified number of errors

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
`-Wremarks` on page 1-61	Issues compiler remarks
`-Wterse` on page 1-61	Issues the briefest form of compiler warning, errors, and remarks
`-w` on page 1-61	Disables all warnings
`-warn-protos` on page 1-62	Issues warnings about functions without prototypes
`-workaround <workaround>` on page 1-62	Enables code generator workaround for specific hardware errata
`-write-files` on page 1-65	Enables compiler I/O redirection
`-write-opts` (on page 1-66)	Passes the user options (but not input filenames) via a temporary file
`-xref` *`filename`* on page 1-66	Outputs cross-reference information to the specified file

Table 1-5. C++ Mode Compiler Switches

Switch Name	Description
`-anach` on page 1-67	Supports some language features (anachronisms) that are prohibited by the C++ standard but still in common use
`-eh` on page 1-68	Enables exception handling
`-no-anach` on page 1-69	Disallows the use of anachronisms that are prohibited by the C++ standard
`-no-demangle` on page 1-69	Prevents filtering of any linker errors through the demangler
`-no-eh` on page 1-69	Disables exception-handling

Table 1-5. C++ Mode Compiler Switches (Cont'd)

Switch Name	Description
`-no-implicit-inclusion` (on page 1-69)	Prevents implicit inclusion of source files as a method of finding definitions of template entities to be instantiated
`-no-rtti` on page 1-69	Disables run-time type information
`-rtti` on page 1-70	Enables run-time type information

C/C++ Mode Selection Switch Descriptions

The following command-line switches provide C/C++ mode selection.

-c89

The `-c89` switch directs the compiler to support programs that conform to the ISO/IEC 9899:1990 standard. For greater conformance to the standard, the following switches should be used: `-alttok`, `-const-read-write`, `no-extra-keywords`, and `-pedantic` (see Table 1-4 on page 1-14).

-c++

The `-c++` (C++ mode) switch directs the compiler to assume that the source file(s) are written in ANSI/ISO standard C++ with Analog Devices language extensions.

All the standard features of C++ are accepted in the default mode except exception handling and run-time type identification because these impose a run-time overhead that is not desirable for all embedded programs. Support for these features can be enabled with the `-eh` and `-rtti` switches (see Table 1-4 on page 1-14).

C/C++ Compiler Common Switch Descriptions

The following command-line switches apply in both C and C++ modes.

sourcefile

The *sourcefile* parameter (or parameters) specifies the name of the file (or files) to be preprocessed, compiled, assembled, and/or linked. A file name can include the drive, directory, file name, and file extension. The ccblkfn compiler uses the file extension to determine the operations to perform. Table 1-2 on page 1-8 lists the permitted extensions and matching compiler operations.

-@ filename

The *-@ filename* (command file) switch directs the compiler to read command-line input from the *filename*. The specified *filename* must contain driver options but may also contain source *filenames* and environment variables. It can be used to store frequently used options as well as to read from a file list.

-A name(tokens)

The -A (assert) switch directs the compiler to assert *name* as a predicate with the specified *tokens*. This has the same effect as the #assert preprocessor directive. The following assertions are predefined.

system	embedded
machine	adspblkfn
cpu	adspblkfn
compiler	ccblkfn

The -A name(value) switch is equivalent to including

```
#assert name(value)
```

in your source file, and both may be tested in a preprocessor condition in the following manner:

```
#if #name(value)
    // do something
#else
    // do something else
#endif
```

For example, the default assertions may be tested as:

```
#if #machine(adspblkfn)
    // do something
#endif
```

The parentheses in the assertion need quotes when using the `-A` switch, to prevent misinterpretation. No quotes are needed for a `#assert` directive in a source file.

-alttok

The `-alttok` (alternative tokens) switch directs the compiler to allow digraph sequences in C and C++ source files. Additionally, the switch enables the recognition of these alternative operator keywords in C++ source files:

Keyword	Equivalent
and	&&
and_eq	&=
bitand	&
bitor	|
compl	~
not	!
not_eq	!=
or	||
or_eq	|=
xor	^
xor_eq	^=

To use alternative tokens in C, you should use `#include <iso646.h>`.

-bss

The `-bss` switch causes the compiler to place global zero-initialized data into a BSS-style section (called "`bsz`"), rather than into the normal global data section. This is the default mode. See also the `-no-bss` switch (on page 1-41).

-build-lib

The `-build-lib` (build library) switch directs the compiler to use the librarian to produce a library file (`.dlb`) instead of using the linker to produce an executable file (`.dxe`). The `-o` option (see on page 1-47) must be used to specify the name of the resulting library.

-C

The `-C` (comments) switch, which is only active in combination with the `-E`, `-EE`, `-ED`, `-P` or `-PP` switches, directs the preprocessor to retain comments in its output.

-c

The `-c` (compile only) switch directs the compiler to compile and/or assemble the source files, but to stop before linking. The output is an object file (`.doj`) for each source file.

-const-read-write

The `-const-read-write` switch directs the compiler to specify that constants may be accessed as read-write data (as in ANSI C). The compiler's default behavior assumes that data referenced through `const` pointers never changes.

The `-const-read-write` switch changes the compiler's behavior to match the ANSI C assumption, which is that other `non-const` pointers may be used to change the data at some point.

-cplbs

The `-cplbs` (CPLBs are active) switch instructs the compiler to assume that all memory accesses will be validated by the Blackfin processor's memory protection hardware. This switch is best used in conjunction with the `-workaround` switches, as it allows the compiler to identify situations where the Cache Protection Lookaside Buffers (CPLBs) will avoid problems, thus avoiding the need for extra workaround instructions.

-Dmacro[=definition]

The `-D` (define macro) switch directs the compiler to define a macro. If you do not include the optional definition string, the compiler defines the macro as the string '`1`'. Note that the compiler processes `-D` switches on the command line before any `-U` (undefine macro) switches.

Invoke this switch with the **Preprocessor definitions** check box located in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **Preprocessor** category.

-debug-types <file.h>

The `-debug-types` switch builds an `*.h` file directly and writes a complete set of debugging information for the header file. The `-g` option (on page 1-33) need not be specified with the `-debug-types` option because it is implied. For example,

```
ccblkfn -debug-types anyHeader.h
```

-decls-<weak|strong>

The `-decls-<weak|strong>` switch controls how the compiler interprets uninitialized global variable definitions, such as `int x;`. The `-decls-strong` switch treats this as equivalent to `int x = 0;`, specifying that other definintions of the same variable in other modules cause a "multiply-defined symbol" error. The `-decls-weak` switch treats this as equivalent to "`extern int x;`", such as a declaration of a symbol that is defined in another module. The default is `-decls-strong`. ANSI C behavior is `-decls-weak`.

-default-linkage-{asm|C|C++}

The `-default-linkage` (default linkage type) switch directs the compiler to set the default linkage type: assembler (`-default-linkage-asm`), C (`-default-linkage-C`) or C++ (`-default-linkage-C++`). C linkage is the default type in C mode, and C++ linkage is the default type in C++ mode.

-double-size-any

The `-double-size-any` switch specifies that the resulting object files should be marked in such a way that will enable them to be linked against objects built with `doubles` either 32-bit or 64-bit in size. Refer to "Using Data Storage Formats" on page 1-261 for more information on data types

-double-size-{32 | 64}

The `-double-size-32` (`double` is 32 bits) and `-double-size-64` (`double` is 64 bits) switches determine the size of the `double` data type. The default is `-double-size-32`, where `double` is a 32-bit data type.
The `-double-size-64` switch promotes `double` to be a 64-bit data type, making it equivalent to `long double`. This switch does not affect the sizes of either `float` or `long double`. Refer to "Using Data Storage Formats" on page 1-261 for more information on data types.

Invoke this switch with the **Double size** radio buttons located in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **Compile** category, **Processor** (1) subcategory.

-dry

The `-dry` (a verbose dry run) switch directs the compiler to display main driver actions, but not to perform them.

-dryrun

The `-dryrun` (a terse dry run) switch directs the compiler to display top-level driver actions, but not to perform them.

-E

The `-E` (stop after preprocessing) switch directs the compiler to stop after the C/C++ preprocessor runs (without compiling). The output (preprocessed source code) prints to the standard output stream unless the output file is specified with the `-O` switch (on page 1-45). Note that the `-C` switch can be used with the `-E` switch.

-ED

The `-ED` (run after preprocessing to file) switch directs the compiler to write the output of the C/C++ preprocessor to a file named "`original_filename.i`". After preprocessing, compilation proceeds normally.

-EE

The `-EE` (run after preprocessing) switch directs the compiler to write the output of the C/C++ preprocessor to standard output. After preprocessing, compilation proceeds normally.

-extra-keywords

The `-extra-keywords` (enable short-form keywords) switch directs the compiler to recognize the Analog Devices keyword extensions to ANSI/ISO standard C/C++ without leading underscores, which can affect conforming ANSI/ISO C/C++ programs. This is the default mode.

The `-no-extra-keywords` switch (on page 1-42) can be used to disallow support for the additional keywords. Table 1-7 on page 1-79 provides a list and a brief description of keyword extensions.

-extra-loop-loads

The `-extra-loop-loads` switch provides the compiler with extra freedom to read more memory locations than required, within a loop, in order to generate the best code. For example, if a loop indicated that the compiler should read elements `arr[0]..arr[59]` and sum them, the `-extra-loop-loads` switch would indicate that the compiler was also allowed to read element `arr[60]`.

-fast-fp

The `-fast-fp` (fast floating point) switch directs the compiler to link with the high-speed floating-point emulation library. This library relaxes some of the IEEE floating-point standard's rules for checking inputs against Not-a-Number (NaN), in the interests of performance. This switch is a default. See also the `-ieee-fp` switch (on page 1-36). Refer to "Using Data Storage Formats" on page 1-261 for more information on data types.

-flags{-asm | -compiler | -lib | -link | -mem} switch [,switch2 [,...]]

The `-flags` (command-line input) switch directs the compiler to pass command-line switches to the other build tools.

These tools are:

Option	Tool
-flags-asm	Assembler
-flags-compiler	Compiler
-flags-lib	Library Builder
-flags-link	Linker
-flags-mem	Memory Initializer

-force-circbuf

The -force-circbuf (circular buffer) switch instructs the compiler to make use of circular buffer facilities, even if the compiler cannot verify that the circular index or pointer is always within the range of the buffer. Without this switch, the compiler's default behavior is conservative, and does not use circular buffers unless it can verify that the circular index or pointer is always within the circular buffer range. See "Circular Buffer Built-In Functions" on page 1-134.

-force-link

The -force-link (force stack frame creation) switch directs the compiler to always create a new stack frame for leaf functions.

This is selected by default if the -g switch (on page 1-33) is selected as it improves the quality of debugging information, but can be switched off with -no-force-link. When -p (on page 1-47) is selected, this switch is always in force. See also -no-force-link switch (on page 1-36).

-fp-associative

The -fp-associative switch directs the compiler to treat floating-point multiplication and addition as associative operations. This switch is on by default.

-full-io

The `-full-io` switch links the application with a third-party, proprietary I/O library. The third-party I/O Library provides a complete implementation of the ANSI C Standard I/O functionality but at the cost of performance (compared to the Analog Devices I/O Library). For more details, see "stdio.h" on page 3-21.

-full-version

The `-full-version` (display version) switch directs the compiler to display version information for all the compilation tools as they process each file.

-g

The `-g` (generate debug information) switch directs the compiler to output symbols and other information used by the debugger.

If the `-g` switch is used with the `-O` (enable optimization) switch, the compiler performs standard optimizations. The compiler also outputs symbols and other information to provide limited source-level debugging through the VisualDSP++ IDDE (debugger). This combination of options provides line debugging and global variable debugging.

Invoke this switch by selecting the **Generate debug information** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **General** category.

When the `-g` and `-O` switches are specified, no debug information is available for local variables and the standard optimizations can sometimes rearrange program code in a way that produces inaccurate line number information. For full debugging capabilities, use the `-g` switch without the `-O` switch. See also the `-Og` switch (on page 1-46).

-guard-vol-loads

The `-guard-vol-loads` (guard volatile loads) switch disables interrupts during volatile loads. A load can be interrupted before completion and restarted once the interrupt completes. If the load is to a device register, this can have undesirable side-effects. The `-guard-vol-loads` switch disables interrupts before issuing a volatile load and re-enables interrupts after the load to avoid this problem.

-H

The `-H` (list headers) switch directs the compiler to output a list of the files included by the preprocessor via the `#include` directive, without compiling. The `-o` switch (on page 1-47) may be used to specify the redirection of the list to a file.

-HH

The `-HH` (list headers and compile) switch directs the compiler to print to the standard output file stream a list of the files included by the preprocessor via the `#include` directive. After preprocessing, compilation proceeds normally.

-h[elp]

The `-help` (command-line help) switch directs the compiler to output a list of command-line switches with a brief syntax description.

-I directory

The `-I` *`directory`* `[{,|;} directory...]` (include search directory) switch directs the C/C++ preprocessor to append the directory (directories) to the search path for `include` files. This option can be specified more than once; all specified directories are added to the search path.

Include files, whose names are not absolute path names and that are enclosed in "..." when included, are searched for in the following directories in this order:

1. The directory containing the current input file (the primary source file or the file containing the `#include`)
2. Any directories specified with the `-I` switch in the order they are listed on the command line
3. Any directories on the standard list:
   `<VisualDSP++ install dir>/.../include`

If a file is included using the <...> form, this file is only searched for by using directories defined in items 2 and 3 above.

-I-

The `-I-` (start include directory list) switch establishes the point in the include directory list at which the search for header files enclosed in angle brackets begins. Normally, for header files enclosed in double quotes, the compiler searches in the directory containing the current input file; then the compiler reverts back to looking in the directories specified with the `-I` switch; and then the compiler searches in the standard `include` directory.

For header files in angle brackets, the compiler performs the latter two searches only.

It is possible to replace the initial search (within the directory containing the current input file) by placing the `-I-` switch at the point on the command line where the search for all types of header file begins. All `include` directories on the command line specified before the `-I-` switch are used only in the search for header files that are enclosed in double quotes.

This switch removes the directory containing the current input file from the `include` directory list.

-i

The `-i` (less includes) switch may be used with the `-H`, `-HH`, `-M`, or `-MM` switches to direct the compiler to only output header details (`-H`, `-HH`) or makefile dependencies (`-M`, `-MM`) for `include` files specified in double quotes.

-ieee-fp

The `-ieee-fp` (slower floating point) switch directs the compiler to link with the fully-compliant floating-point emulation library. This library obeys all the IEEE floating-point standard's rules, and incurs a performance penalty when compared with the default floating-point emulation library. See also the `-fast-fp` switch (on page 1-31). Refer to "Using Data Storage Formats" on page 1-261 for more information on data types.

-include filename

The `-include` *`filename`* (include file) switch directs the preprocessor to process the specified file before processing the regular input file. Any `-D` and `-U` options on the command line are processed before an `-include` file.

-ipa

The `-ipa` (interprocedural analysis) switch turns on Interprocedural Analysis (IPA) in the compiler. This option enables optimization across the entire program, including between source files that were compiled separately. The `-ipa` option should be applied to all C and C++ files in the program. For more information, see "Interprocedural Analysis" on page 1-74. Specifying `-ipa` also implies setting the `-O` switch (on page 1-45).

Invoke this switch by selecting the **Interprocedural optimization** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **General** category.

-jcs2l

This switch requests the linker to convert compiler-generated short jumps to long jumps when necessary.

-jcs2l+

The `-jcs2l+` switch requests the linker to convert compiler-generated short jumps to long jumps when necessary, but uses the P1 register for indirect jumps/calls when long jumps/calls are insufficient. Enabled by default.

-jump-<constdata|data|code>

The `-jump-<constdata|data|code>` switch determines which section the compiler uses to store jump tables that are generated for some C "switch" statements. The `-jump-constdata` option is the default for this switch; it causes jump tables to be placed in the `constdata` section as read-only data. The `-jump-data` option places the jump tables into `data1`, which is the default for normal read-write data (which also used to be the compiler's default). The `-jump-code` option causes the compiler to place the jump table into the `program` section.

Since L1 Instruction memory is not readable as data, you may not use `-jump-code` on programs that may map into L1 instruction memory, as an invalid-address exception would be raised when the program attempts to read the table.

-L directory[{,|;} **directory…**]

The `-L` *`directory`* (library search directory) switch directs the linker to append the directory (or directories) to the search path for library files.

-l library

The `-l` (link library) switch directs the linker to search the library for functions and global variables when linking. The library name is the portion of the file name between the `lib` prefix and `.dlb` extension. For example, the `-lc` compiler switch directs the linker to search in the library named `c`. This library resides in a file named `libc.dlb`.

All object files should be listed on the command line before listing libraries using the `-l` switch. When a reference to a symbol is made, the symbol definition will be taken from the left-most object or library on the command line that contains the global definition of that symbol. If two objects on the command line contain definitions of the symbol `x`, `x` will be taken from the left-most object on the command line that contains a global definition of `x`.

If one of the definitions for `x` comes from user objects, and the other from a user library, and the library definition should be overridden by the user object definition, it is important that the user object comes before the library on the command line.

Libraries included in the default `.LDF` file are searched last for symbol definitions.

-M

The `-M` (generate make rules only) switch directs the compiler not to compile the source file, but to output a rule suitable for the make utility, describing the dependencies of the main program file.

The format of the make rule output by the preprocessor is:

```
object-file: include-file ...
```

-MD

The `-MD` (generate make rules and compile) switch directs the preprocessor to print to a file called `original_filename.d` a rule describing the dependencies of the main program file. After preprocessing, compilation proceeds normally. See also the `-Mo` switch.

-MM

The `-MM` (generate make rules and compile) switch directs the preprocessor to print to the standard output stream a rule describing the dependencies of the main program file. After preprocessing, compilation proceeds normally.

-Mo filename_

The `-Mo` *`filename`* (preprocessor output file) switch directs the compiler to use *`filename`* for the output of `-MD` or `-ED` switches.

-Mt filename

The `-Mt` *`filename`* (output make rule for the named source) switch modifies the target of dependencies, renaming the target to filename. It only has an effect when used in conjunction with the `-M` or `-MM` switch.

-MQ

The `-MQ` switch directs the compiler not to compile the source file but to output a rule. In addition, the `-MQ` switch does not produce any notification when input files are missing.

-map filename

The `-map filename` (generate a memory map) switch directs the linker to output a memory map of all symbols. The map file name corresponds to the *`filename`* argument. For example, if the file name argument is `test`, the map file name is `test.xml`. The `.xml` extension is added where necessary.

-mem

The `-mem` (invoke memory initializer) switch causes the compiler to invoke the Memory Initializer tool after linking the executable file. The MemInit tool can be controlled through the `-flags-mem` switch (on page 1-31).

-multicore

The `-multicore` switch indicates to the compiler that the application is being built for use in a dual-core environment, such as the ADSP-BF561 Blackfin processor. It indicates that both cores are operating at once, and therefore the application is linked against versions of the libraries that including locking and per-core private storage. The `-multicore` switch defines the `__ADI_MULTICORE` macro to the value "1" at both compile-time and link-time.

The `-multicore` switch is not supported in conjunction with the `-pguide`, `-ipa`, `-p`, `-p1` or `-p2` switches.

-multiline

The `-multiline` switch enables a compiler GNU compatibility mode which allows string literals to span multiple lines without the need for a "\" at the end of each line. This is the default mode.

-no-alttok

The `-no-alttok` (disable alternative tokens) switch directs the compiler not to accept alternative operator keywords and digraph sequences in the source files. This is the default mode. For more information, see "-alttok" on page 1-26.

-no-annotate

The `-no-annotate` (disable assembly annotations) directs the compiler not to annotate assembly files. The default behavior is that whenever optimizations are enabled all assembly files generated by the compiler are annotated with information on the performance of the generated assembly. See "Assembly Optimizer Annotations" on page 2-54 for more details on this feature.

-no-bss

The `-no-bss` switch causes the compiler to keep zero-initialized and non-zero-initialized data in the same data section, rather than separating zero-initialized data into a different, BSS-style section. See also the `-bss` switch (on page 1-27).

-no-builtin

The `-no-builtin` (no built-in functions) switch directs the compiler to ignore any built-in functions that do not begin with two underscores (__). Note that this switch influences many functions. This switch also predefines the `__NO_BUILTIN` preprocessor macro. For more information, see "Compiler Built-In Functions" on page 1-103.

Invoke this switch by selecting the **Disable builtin functions** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **General** category.

-no-circbuf

The `-no-circbuf` (no circular buffer) switch directs the compiler not to automatically use circular buffer mechanisms (such as for referencing `array[i % n]`). Use of the `circindex()` and `circptr()` functions (that is, explicit circular buffer operations) is not affected.

-no-defs

The `-no-defs` (disable defaults) switch directs the compiler not to define any default preprocessor macros, include directories, library directories, libraries, or run-time headers.

-no-extra-keywords

The `-no-extra-keywords` (disable short-form keywords) switch directs the compiler not to recognize the Analog Devices keyword extensions that might affect conformance to ANSI/ISO standards for the C and C++ languages. Keywords, such as `asm`, may be used as identifiers in conforming programs. Alternate keywords, which are prefixed with two leading underscores, such as `__asm`, continue to work.

-no-force-link

The `-no-force-link` (do not force stack frame creation) switch directs the compiler not to create a new stack frame for leaf functions.

This switch is most useful in combination with the `-g` switch (on page 1-33) when debugging optimized code. When optimization is requested, the compiler does not generate stack frames for functions that do not need them, which improves both the size and speed of the code but reduces the quality of information displayed in the debugger. Therefore, when the `-g` switch is used, the compiler by default always generates a stack frame. Consequently, the code generated with the `-g` switch is differ-

ent from the code generated without using this switch which may result in different incorrect program behavior. The `-no-force-link` switch causes the same code to be generated whether `-g` is used or not.

-no-fp-associative

The `-no-fp-associative` switch directs the compiler NOT to treat floating-point multiplication and addition as associative operations.

-no-full-io

The `-no-full-io` switch links the application with the Analog Devices I/O library which contains a faster implementation of C Standard I/O than the alternative third-party I/O Library (see "-full-io" on page 1-33). The functionality provided by the Analog Devices I/O library is not as comprehensive as the third-party I/O library. For more details, "stdio.h" on page 3-21.

This switch passes the macro `_ADI_LIBIO` to the compiler and linker. This switch is enabled by default.

-no-int-to-fract

The `-no-int-to-fract` (disable conversion of integer to fractional arithmetic) switch directs the compiler not to turn integer arithmetic into fractional arithmetic. For example, a statement such as

```
short a = ((b*c)>>15);
```

may be changed, by default, into a fractional multiplication. The saturation properties of integer and fractional arithmetic are different; therefore, if the expression overflows, then the results differ. Specifying the `-no-int-to-fract` switch disables this optimization.

-no-jcs2l

The `-no-jcs2l` switch prevents the linker from converting compiler-generated `short` jumps to `long` jumps.

-no-jcs2l+

The `-no-jcs2l+` switch prevents the linker from converting compiler-generated `short` jumps to `long` jumps using register `P1`.

-no-mem

The `-no-mem` switch causes the compiler not to invoke the Memory Initializer tool after linking the executable. This is the default setting. See also "-mem" on page 1-40.

-no-multiline

The `-no-multiline` switch disables a compiler GNU compatibility mode which allows string literals to span multiple lines without the need for a "\" at the end of each line.

-no-saturation

The `-no-saturation` switch directs the compiler not to introduce faster operations in cases where the faster operation would saturate (if the expression overflowed) when the original operation would have wrapped the result.

-no-std-ass

The `-no-std-ass` (disable standard assertions) switch prevents the compiler from defining the standard assertions. See the `-A` switch (on page 1-25) for the list of standard assertions.

-no-std-def

The `-no-std-def` (disable standard macro definitions) switch prevents the compiler from defining default preprocessor macro definitions.

-no-std-inc

The `-no-std-inc` (disable standard include search) switch directs the C/C++ preprocessor to search only for header files in the current directory and directories specified with the `-I` switch.

Invoke this switch by selecting the **Ignore standard include paths** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **Preprocessor** category.

-no-std-lib

The `-no-std-lib` (disable standard library search) switch directs the linker to limit its search for libraries to directories specified with the `-L` switch (on page 1-37). The compiler also defines `__NO_STD_LIB` during the linking stage and passes it to the linker, so that the `SEARCH_DIR` directives in the `.LDF` file can de disabled.

-no-threads

The `-no-threads` (disable thread-safe build) switch specifies that all compiled code and libraries used in the build need not be thread safe. This is the default setting when the `-threads` (enable thread-safe build) switch is not used.

-O[0|1]

The `-O` (enable optimizations) switch directs the compiler to produce code that is optimized for performance. Optimizations are not enabled by default for the compiler. The switch setting `-O` or `-O1` turns optimization on, while setting `-O0` turns off all optimizations.

Invoke this switch by selecting the **Enable optimization** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **General** category.

-Oa

The `-Oa` (automatic function inlining) switch enables the inline expansion of C/C++ functions, which are not necessarily declared inline in the source code. The amount of auto-inlining the compiler performs is controlled using the `-Ov` (optimize for speed versus size) switch (on page 1-47). Therefore, use of `-Ov100` indicates that as many functions as possible will be auto-inlined, whereas `-Ov0` prevents any function from being auto-inlined. Specifying `-Oa` also implies the use of `-O`.

-Ofp

The `-Ofp` (frame pointer optimization) switch directs the compiler to offset the Frame Pointer within a function, this allows the compiler to use more short load and store instructions. Specifying `-Ofp` also implies the use of `-O`.

Specifying this switch reduces the capabilities of the debugger for source-level debugging actions when used with `-g`, since the active call frames cannot be followed beyond an active function with a Frame Pointer offset. Debugger facilities affected by the `-Ofp` switch are: Call Stack, Step Over, and Step Out Of.

-Og

The `-Og` switch enables a compiler mode that attempts to perform optimizations while still preserving the debugging information. It is meant as an alternative for those users who want a debuggable program but who are also concerned about the performance of their debuggable code.

-Os

The `-Os` (enable code size optimization) switch directs the compiler to produce code that is optimized for size. This is achieved by performing all optimizations except those that increase code size. The optimizations not performed include loop unrolling and jump avoidance.

-Ov num

The `-Ov num` (optimize for speed versus size) switch directs the compiler to produce code that is optimized for speed versus size. The "`num`" should be an integer between 0 (purely size) and 100 (purely speed).

-o filename

The `-o filename` (output file) switch directs the compiler to use *`filename`* for the name of the final output file.

-P

The `-P` (omit line numbers) switch directs the compiler to stop after the C/C++ preprocessor runs (without compiling) and to omit the `#line` preprocessor directives (with line number information) in the output from the preprocessor. The `-C` switch can be used with the `-P` switch to retain comments.

-PP

The `-PP` (omit line numbers and compile) switch is similar to the `-P` switch; however, it does not halt compilation after preprocessing.

-p[1|2]

The `-p` (generate profiling implementation) switch directs the compiler to generate the additional instructions needed to profile the program by recording the number of cycles spent in each function.

The `-p1` switch causes the program being profiled to write the information to a file called `mon.out`. The `-p2` switch changes this behavior to write the information to the standard output file stream. The `-p` switch writes the data to both `mon.out` and the standard output stream. For more information on profiling, see "Profiling With Instrumented Code" on page 1-197.

-path {-asm | -compiler | -lib | -link | -mem} filename

The `-path` (tool location) switch directs the compiler to use the specified component in place of the default-installed version of the compilation tool. The component comprises a relative or absolute path to its location. Respectively, the tools are the assembler, compiler, librarian, linker or memory initializer. Use this switch when overriding the normal version of one or more of the tools. The `-path` switch also overrides the directory specified by the `-path-install` switch (on page 1-48).

-path-install directory

The `-path-install` *`directory`* (installation location) switch directs the compiler to use the specified directory as the location for all compilation tools instead of the default path. This is useful when working with multiple versions of the tool set.

You can selectively override this switch with the `-path-`*`tool`* switch.

-path-output directory

The `-path-output` *`directory`* (non-temporary files location) switch directs the compiler to place output files in the specified directory.

-path-temp directory

The `-path-temp` *`directory`* (temporary files location) switch directs the compiler to place temporary files in the specified directory.

-pch

The `-pch` (recompiled header) switch directs the compiler to automatically generate and use precompiled header files. A precompiled output header has a `.pch` extension attached to the source file name. By default, all precompiled headers are stored in a directory called `PCHRepository`.

Precompiled header files can significantly speed compilation; precompiled headers tend to occupy more disk space.

-pchdir directory

The `-pchdir` *`directory`* (locate precompiled header repository) switch specifies the location of an alternative directory for storing and invocation of precompiled header files. If the directory does not exist, the compiler creates it. Note that the `-o` (output) switch does not influence the `-pchdir` option.

-pedantic

The `-pedantic` (ANSI standard warning) switch causes the compiler to issue a warning for each construct found in your program that does not strictly conform to ANSI/ISO standard C or C++ language.

The compiler may not detect all such constructs. In particular, the `-pedantic` switch does not cause the compiler to issue errors when Analog Devices keyword extensions are used.

-pedantic-errors

The `-pedantic-errors` (ANSI standard errors) switch causes the compiler to issue an error instead of a warning for cases described in the `-pedantic` switch.

-pguide

The `-pguide` (PGO) switch causes the compiler to add instrumentation for the gathering of a profile as the first stage of performing profile-guided optimization.

-pplist filename

The `-pplist` *`filename`* (preprocessor listing) directs the preprocessor to output a listing to the named file. When more than one source file is preprocessed, the listing file contains information about the last file processed. The generated file contains raw source lines, information on transitions into and out of include files, and diagnostics generated by the compiler.

Each listing line begins with a key character that identifies its type, such as

Character	Meaning
`N`	Normal line of source
`X`	Expanded line of source
`S`	Line of source skipped by `#if` or `#ifdef`
`L`	Change in source position
`R`	Diagnostic message (remark)
`W`	Diagnostic message (warning)
`E`	Diagnostic message (error)
`C`	Diagnostic message (catastrophic error)

-proc processor

The `-proc` (target processor) switch specifies the compiler produces code suitable for the specified processor. Refer to "Supported Processors" for the list of supported Blackfin processors. For example,

```
ccblkfn -proc ADSP-BF535 -o bin\p1.doj p1.asm
```

If no target is specified with the `-proc` switch, the system uses the `ADSP-BF532` setting as a default.

When compiling with the `-proc` switch, the appropriate processor macro and `__ADSPBLACKFIN__` preprocessor macro are defined as 1. When the target processor is a ADSP-BF531, ADSP-BF532, ADSP-BF533, ADSP-BF534, ADSP-BF536, ADSP-BF537, ADSP-BF538, ADSP-BF539, ADSP-BF561 or ADSP-BF566 processor, the compiler additionally defines macro `__ADSPLPBLACKFIN__` to 1. For example, when `-proc ADSP-BF531` is used, macros `__ADSPBF531__`, `__ADSPBLACKFIN__`, and `__ADSPLPBLACKFIN__` are all pre-defined to 1 by the compiler.

See also "-si-revision version" on page 1-55 for more information on the silicon revision of the specified processor.

-progress-rep-func

The `-progress-rep-func` switch provides feedback on the compiler's progress that may be useful when compiling and optimizing very large source files. It issues a "warning" message each time the compiler starts compiling a new function. The "warning" message is a remark that is disabled by default, and this switch enables the remark as a warning. The switch is equivalent to `-Wwarn=cc1472`.

-progress-rep-gen-opt

The `-progress-rep-gen-opt` switch provides feedback on the compiler's progress that may be useful when compiling and optimizing a very large, complex function. It issues a "warning" message each time the compiler starts a new generic optimization pass on the current function. The "warning" message is a remark that is disabled by default, and this switch enables the remark as a warning. The switch is equivalent to `-Wwarn=cc1473`.

-progress-rep-mc-opt

The `-progress-rep-mc-opt` switch provides feedback on the compiler's progress that may be useful when compiling and optimizing a very large, complex function. It issues a "warning" message each time the compiler starts a new machine-specific optimization pass on the current function. The "warning" message is a remark that is disabled by default, and this switch enables the remark as a warning. The switch is equivalent to `-Wwarn=cc1474`.

-R directory[{;|,}**directory …**]

The `-R` *`directory`* (add source directory) switch directs the compiler to add the specified directory to the list of directories searched for source files.

On Windows platforms, multiple source directories are given as a comma- or semicolon-separated list. The compiler searches for the source files in the order specified on the command line. The compiler searches the specified directories before reverting to the current directory. This switch is dependent on its position on the command line; that is, it effects only source files that follow it.

Source files, whose file names begin with `/`, `./` or `../` (or Windows equivalent) or contain drive specifiers (on Windows platforms), are not affected by this option.

-R-

The `-R-` (disable source path) switch removes all directories from the standard search path for source files, effectively disabling this feature.

This option is position-dependent on the command line; it only affects files following it.

-reserve register[, register ...]

The `-reserve` (reserve register) switch directs the compiler not to use the specified registers. Only the `m3` register can be reserved.

-S

The `-S` (stop after compilation) switch directs the compiler to stop compilation before running the assembler. The compiler outputs an assembly file with an `.s` extension.

Invoke this switch by selecting the **Stop after: Compiler** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **General** category.

-s

The `-s` (strip debug information) switch directs the compiler to remove debug information (symbol table and other items) from the output executable file during linking.

-sat32

The `-sat32` (32-bit saturation) switch directs the compiler to generate code to saturate multiply and accumulate results in the accumulator at 32 bits. This is the default setting.

-sat40

The `-sat40` (40-bit saturation) switch directs the compiler to generate code to saturate multiply and accumulate results in the accumulator at 40 bits, rather than at the default, which saturates at 32 bits.

When compiler optimizations are not enabled, the `-sat40` driver option has minimal impact. The compiler optimizer contains an algorithm to identify a sequence of instructions which could be replaced by a single MAC operation performed using an accumulator. This may result in uses

of the accumulator that appear in optimized code, where they do not in non-optimized code. As a result of this, the switch has a more significant impact when compiler optimizations are enabled.

-save-temps

The `-save-temps` (save intermediate files) switch directs the compiler to retain intermediate files generated and normally removed as part of the various compilation stages. These intermediate files are placed in the `-path-output` specified output directory or the build directory if the `-path-output` switch (on page 1-48) is not used. See Table 1-2 on page 1-8 for a list of intermediate files.

-sdram

The `-sdram` (SDRAM is active) switch instructs the compiler to assume that at least Bank 0 of external SDRAM—the lower 32 Mbytes of space—is active and enabled. This switch is best used in conjunction with the `-workaround` switches (on page 1-62, as it allows the compiler to identify where memory accesses will harmlessly read SDRAM, rather than requiring extra workaround instructions to avoid speculative reads that might cause hardware errors.

-section <id>=<section_name>[,<id>=section_name>]

The `-section` switch controls the placement of types of data produced by the compiler. The data is placed into the section "`section_name`" as provided on the command line.

The compiler currently supports the following section identifier:

`vtbl` – controls placement of the C++ virtual lookup tables

Make sure that the section selected via the command line exists within the `.LDF` file (refer to the *Linker* chapter in the *VisualDSP++ 4.0 Linker and Utilities Manual*).

-show

The `-show` (display command line) switch directs the compiler to display the command-line arguments passed to the driver, including expanded option files and environment variables. This option allows you to ensure that command-line options have been successfully invoked by the driver.

-signed-bitfield

The `-signed-bitfield` (make plain bitfields signed) switch directs the compiler to make bitfields (which have not been declared with an explicit signed or unsigned keyword) signed. This switch does not effect plain one-bit bitfields, which are always unsigned. This is the default mode. See also the `-unsigned-bitfield` switch (on page 1-59).

-signed-char

The `-signed-char` (make char signed) switch directs the compiler to make the default type for `char` signed. The compiler also defines the `__SIGNED_CHARS__` macro. This is the default mode when the `-unsigned-char` switch is not used (on page 1-60).

-si-revision *version*

The `-si-revision` *`version`* (silicon revision) switch directs the compiler to build for a specific hardware revision. Any errata workarounds available for the targeted silicon revision will be enabled. The parameter "*`version`*" represents a silicon revision for the processor specified by the `-proc` switch (on page 1-50).

For example,

```
ccblkfn -proc ADSP-BF535 -si-revision 0.1 proc.c
```

If silicon version "`none`" is used, then no errata workarounds are enabled, whereas specifying silicon version "`any`" will enable all errata workarounds for the target processor.

If the `-si-revision` switch is not used, the compiler will build for the latest known silicon revision for the target processor and any errata workarounds which are appropriate for the latest silicon revision will be enabled.

The directory `Blackfin\lib` contains two sets of libraries: one set (suffixed with "`y`", for example, `libc532y.dlb`) contains workarounds for all known errata in all silicon revisions, the other set is built without any errata workarounds. Within the `lib` sub-directory, there are library directories for each silicon revision; these libraries have been built with errata workarounds appropriate for the silicon revision enabled. Note that an individual set of libraries may cover more than one specific silicon revision, so if several silicon revisions are affected by the same errata, then one common set of libraries might be used.

The `__SILICON_REVISION__` macro is set by the compiler to two hexadecimal digits representing the major and minor numbers in the silicon revision. For example, `1.0` becomes `0x100` and `10.21` becomes `0xa15`.

If the silicon revision is set to "`any`", the `__SILICON_REVISION__` macro is set to `0xffff` and if the `-si-revision` switch is set to "`none`", the compiler will not set the `__SILICON_REVISION__` macro.

The compiler driver will pass the `-si-revision <silicon version>` switch when invoking another VisualDSP++ tool, for example when the compiler driver invokes the assembler and linker.

Use `http://www.analog.com/processors/technicalSupport/hardwareAnomalies.html` to get more information on specific anomalies (including anomaly IDs).

-structs-do-not-overlap

The `-structs-do-not-overlap` switch specifies that the source code being compiled contains no structure copies such that the source and the destination memory regions overlap each other in a non-trivial way.

For example, in the statement

```
*p = *q;
```

where p and q are pointers to some structure type S, the compiler, by default, always ensures that, after the assignment, the structure pointed to by "p" contains an image of the structure pointed to by "q" prior to the assignment. In the case where p and q are not identical (in which case the assignment is trivial) but the structures pointed to by the two pointers may overlap each other, doing this means that the compiler must use the functionality of the C library function "memmove" rather than "memcpy".

It is slower to use "memmove" to copy data than it is to use "memcpy". Therefore, if your source code does not contain such overlapping structure copies, you can obtain higher performance by using the command-line switch `-structs-do-not-overlap`.

Invoke this switch from a check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **General** category.

-syntax-only

The `-syntax-only` (only check syntax) switch directs the compiler to check the source code for syntax errors and warnings. No output files are generated with this switch.

-sysdefs

The `-sysdefs` (system definitions) switch directs the compiler to define several preprocessor macros describing the current user and user's system. The macros are defined as character string constants and are used in functions with null-terminated string arguments.

These macros are defined if the system returns information for them.

Note that the __MACHINE__, __GROUPNAME__, and __REALNAME__ macros are not available on Windows platforms.

Macro	Description
__HOSTNAME__	The name of the host machine
__MACHINE__	The machine type of the host machine
__SYSTEM__	The Operating System name of the host machine
__USERNAME__	The current user's login name
__GROUPNAME__	The current user's group name
__REALNAME__	The current user's real name.

-T filename

The `-T filename` (Linker Description File) switch directs the linker to use the specified Linker Description File (`.LDF`) as control input for linking. If `-T` is not specified, a default `.LDF` file is selected based on the processor variant.

-threads

When used, the `-threads` switch defines the macro `__ADI_THREADS` as one (1) at the compile, assemble and link phases of a build. This specifies that certain aspects of the build are to be done in a thread-safe way.

The use of thread-safe libraries is necessary in conjunction with the `-threads` flag when using the VisualDSP++ Kernel (VDK). The thread-safe libraries can be used with other RTOSs but this requires the definition of various VDK interfaces. The use of `-threads` switch does not imply that the compiler will produce thread-safe code when compiling C/C++ source. It is the user's responsibility that multi-threaded programming practises are employed in their code (such as using semaphores to access shared data).

-time

The `-time` (tell time) switch directs the compiler to display elapsed time as part of the output information on each part of the compilation process.

-U macro

The `-U` (undefine macro) switch directs the compiler to undefine macros. If you specify a macro name, it is undefined. Note the compiler processes all `-D` (define macro) switches on the command line before any `-U` (undefine macro) switches.

Invoke this switch by selecting the **Preprocessor undefines** field in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **Preprocessor** category.

-unsigned-bitfield

The `-unsigned-bitfield` (make plain bitfields unsigned) switch directs the compiler to make bitfields (which have not been declared with an explicit signed or unsigned keyword) unsigned. This switch does not effect plain one-bit bitfields, which are always unsigned.

For example, given the declaration

```
struct {
   int a:2;
   int b:1;
   signed int c:2;
   unsigned int d:2;
} x;
```

the `bitfield` values are:

Field	-signed-bitfield	-unsigned-bitfield	Why
x.a	-2..1	0..3	Plain field
x.b	0..1	0..1	One bit

x.c	-2..1	-2..1	Explicit signed
x.d	0..3	0..3	Explicit unsigned

See also the `-signed-bitfields` switch (on page 1-55).

-unsigned-char

The `-unsigned-char` (make char unsigned) switch directs the compiler to make the default type for `char` unsigned. The compiler also undefines the `__SIGNED_CHARS__` preprocessor macro.

-v

The `-v` (version and verbose) switch directs the compiler to display the version and command-line information for all the compilation tools as they process each file.

-verbose

The `-verbose` (display command line) switch directs the compiler to display command-line information for all the compilation tools as they process each file.

-version

The `-version` (display version) switch directs the compiler to display its version information.

-W{error|remark|suppress|warn} number[, number...]

The `-W {...}` *`number`* (override error message) switch directs the compiler to override the severity of the specified diagnostic messages (errors, remarks, or warnings). The *`number`* argument specifies the message to override.

At compilation time, the compiler produces a number for each specific compiler diagnostic message. The `{D}` (discretionary) suffix after the diagnostic message number indicates that the diagnostic may have its severity overridden. Each diagnostic message is identified by a number that is used across all compiler software releases.

-Werror-limit number

The `-Werror-limit` (maximum compiler errors) switch lets you set a maximum number of errors for the compiler before it aborts.

-Wremarks

The `-Wremarks` (enable diagnostic warnings) switch directs the compiler to issue remarks, which are diagnostic messages that are milder than warnings.

You can invoke this switch by selecting the **Enable remarks** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **Warning** selection.

-Wterse

The `-Wterse` (enable terse warnings) switch directs the compiler to issue the briefest form of warnings. This also applies to errors and remarks.

-w

The `-w` (disable all warnings) switch directs the compiler not to issue warnings.

Invoke this switch by selecting the **Disable all warnings and remarks** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **Warning** selection.

-warn-protos

The `-warn-protos` (warn if incomplete prototype) switch directs the compiler to issue a warning when it calls a function for which an incomplete function prototype has been supplied. This option has no effect in C++ mode.

-workaround <workaroundid>[,<workaroundid> …]

The `-workaround <workaroundid>[,<workaroundid> …]` switch enables compiler code generator workarounds for specific hardware errata.

Use `http://www.analog.com/processors/technicalSupport/hardwareAnomalies.html` to get more information on specific anomalies (including anomaly IDs).

Current valid workarounds are listed in Table 1-6 (workarounds may change).

Table 1-6. Current Workaround Descriptions

WorkaroundID	Workaround Function
all	Indicates that the compiler enables all known workarounds. In addition to the compiler generating errata safe code, the run-time libraries linked using the default `.LDF` files are safe for all workarounds. These alternate form of libraries and object files include a "`y`" suffix in their filenames (for example, `libc535y.dlb`).
avoid-dag1	Instructs the compiler not to use `DAG1` when issuing memory accesses. This avoids anomaly 05-00-0104. The compiler also defines macros `__WORKAROUND_AVOID_DAG1` and `__WORKAROUNDS_ENABLED` at the source, assembly and link build stages when this workaround is enabled.

Table 1-6. Current Workaround Descriptions (Cont'd)

WorkaroundID	Workaround Function
avoid-dag-load-reuse	Instructs the compiler to avoid anomaly 05-00-0087, where a memory access may use the wrong address. The sequence may include: 1. DAGREGx = [memory] 2. some memory or MMR read that stalls 3. any access to memory using DAGREGx within the next four instructions The final memory access may use an incorrect address. This workaround instructs the compiler to avoid this sequence by inserting a NOP between the first load of DAGREGx and the second memory read that may stall. The compiler also defines macros __WORKAROUND_AVOID_DAG_LOAD_REUSE and __WORKAROUNDS_ENABLED at the source, assembly and link build stages, when this workaround is enabled.
csync	The csync workaround ensures that the compiler will avoid anomaly 05-00-0048, where the processor speculatively executes a memory access to an address as part of an instruction that may not be committed. If the address is invalid, exceptions may be raised. The compiler avoids the anomaly by automatically inserting CSYNC instructions after conditional branches when: • the branch is not predicted to be taken • the first instruction following the branch includes a load • the load is not through the stack or frame pointer If any of these conditions are not met, the compiler does not need to insert a CSYNC instruction. However, there are cases where the compiler inserts a CSYNC that is not necessary. These cases are when the pointer always points to a valid address, even if the branch is taken. The compiler also inserts CSYNC instructions before writes to either of the CYCLES or CYCLES2 registers, in the shadow of conditional branches, and after RTI, RTX or RTN instructions. When optimizing for speed rather than space, the compiler converts the inserted CSYNCs into a sequence of NOP instructions; these instructions require more space than the single CSYNC, but do not flush the processor's pipeline. When optimizing for space, the CSYNC instructions are not replaced by NOP sequences. The compiler also defines macros __WORKAROUND_CSYNC and __WORKAROUNDS_ENABLED at the source, assembly and link build stages when this workaround is enabled.
isr-imask-check	Indicates that when compiling interrupt service routines, the compiler should check that IMASK has not been cleared before servicing the interrupt. This avoids anomaly 05-00-0071. The compiler also defines macros __WORKAROUND_IMASK_CHECK and __WORKAROUNDS_ENABLED at the source, assembly and link build stages when this workaround is enabled.

Table 1-6. Current Workaround Descriptions (Cont'd)

WorkaroundID	Workaround Function
isr-ssync	Indicates that when compiling interrupt service routines, the compiler should insert a SSYNC instruction at the start of the ISR, so that context saves do not interfere with cache write-back stores. This avoids anomaly 05-00-0054. The compiler also defines macros __WORKAROUND_SSYNC and __WORKAROUNDS_ENABLED at the source, assembly and link build stages when this workaround is enabled.
wb-dcache	Enables workaround for the anomaly 05-00-0165 "Data cache dirty bit set when a load-miss-fill is in progress". Enabling the workaround has the following effects: • A SSYNC instruction is placed at the beginning of ISRs. • Two NOP instructions are placed after RTI, RTX, RTN and RTE instructions. • A NOP instruction is inserted between conditional branches and load or store instructions. • Because some assembly files may be specified on the command line, the assembler is passed the switch -wb_fix, which results in a NOP being inserted between conditional branches and load or store instructions. When this workaround is enabled, the following macros are defined at the compile, link, and assembly phases: _WORKAROUND_WB_DCACHE and __WORKAROUNDS_ENABLED
wt-dcache	Enables workaround for the anomaly 05-00-0164 "Store to Load Forwarding in Write Through Mode". This workaround includes the wb-dcache workaround, and in addition: • A SSYNC instruction is placed before the RTI instruction in ISRs. • A SSYNC instruction is inserted between the target of a conditional branch and the first store in that basic block. • Because some assembly files may be specified on the command line, the assembler is passed the -wb_wt_fix switch and an SSYNC instruction is inserted between the target of a conditional branch and the first store in that basic block.When this workaround is enabled, the following macros are defined at the compile, link, and assembly phases: __WORKAROUND_WT_DCACHE __WORKAROUNDS_ENABLED
signbits	Indicates that the compiler should avoid generating code where the input register of a signbits instruction is loaded in the immediately preceding instruction. This avoids anomaly 05-00-0127. The compiler also defines macros __WORKAROUND_SIGNBITS and __WORKAROUNDS_ENABLED at the source, assembly and link build stages when this workaround is enabled.

Table 1-6. Current Workaround Descriptions (Cont'd)

WorkaroundID	Workaround Function
killed-mmr-write	Indicates that the compiler should insert a dummy 32-bit system MMR read at the start of ISR code and to avoid 32-bit system MMR write accesses in the 3 slots after a conditional branch which is predicted not taken. This avoids anomaly 05-00-0157. The compiler also defines macros `__WORKAROUND_KILLED_MMR_WRITE` and `__WORKAROUNDS_ENABLED` at the source, assembly and link build stages when this workaround is enabled.

astat-rnd_mod	Enables workaround for the anomaly 05-00-0195 "Latency in Writes to `RND_MOD` bit". Enabling the workaround inserts two `NOP` instructions after any direct write to the `ASTAT` register.
	The compiler generates a write to the `ASTAT` register only when `__builtin_sysreg_write()` is used to write to the `ASTAT` register.
	When this workaround is enabled, the following macros are defined at the compile, link, and assembly phases:
	`__WORKAROUND_ASTAT_RND_MOD`
	`__WORKAROUNDS_ENABLED`

sdram-mmr-read	Enables workaround for the anomaly 05-00-0198 "Memory Access Pipeline Bug". The problem is seen in a sequence of SDRAM load directly followed by an MMR load and will result in an incorrect value being loaded from the MMR.
	When this workaround is enabled, the compiler inserts a `NOP` between the two loads. It also defines the following macros at the compile, link, and assembly phases:
	`__WORKAROUND_SDRAM_MMR_READ`
	`__WORKAROUNDS_ENABLED`

-write-files

The `-write-files` (enable driver I/O redirection) switch directs the compiler driver to redirect the file name portions of its command line through a temporary file. This technique helps to handle long file names, which can make the compiler driver's command line too long for some operating systems.

-write-opts

The `-write-opts` (user options) switch directs the compiler to pass the user options (but not the input file names) to the main driver via a temporary file which can help if the resulting main driver command line is too long.

-xref filename

The `-xref` *`filename`* (cross-reference list) switch directs the compiler to write cross-reference listing information to the specified file. When more than one source file has been compiled, the listing contains information about the last file processed. For each reference to a symbol in the source program, a line of the form

```
symbol-id name ref-code filename line-number column-number
```

is written to the named file. The `symbol-id` represents a unique decimal number for the symbol, and `ref-code` is one of the following characters.

Character	Meaning
`D`	Definition
`d`	Declaration
`M`	Modification
`A`	Address taken
`U`	Used
`C`	Changed (used and modified)
`R`	Any other type of reference
`E`	Error (unknown type of reference)

C++ Mode Compiler Switch Descriptions

The following switches apply only to the C++ compiler.

-anach

The `-anach` (enable C++ anachronisms) directs the compiler to accept some language features that are prohibited by the C++ standard but are still in common use. This is the default mode. Use the `-no-anach` switch for greater standard compliance.

The following anachronisms are accepted in the default C++ mode:

- Overload is allowed in function declarations. It is accepted and ignored.
- Definitions are not required for static data members that can be initialized using default initialization. The anachronism does not apply to static data members of template classes; they must always be defined.
- The number of elements in an array may be specified in an array delete operation. The value is ignored.
- A single `operator++()` and `operator--()` function can be used to overload both prefix and postfix operations.
- The base class name may be omitted in a base class initializer if there is only one immediate base class.
- Assignment to `this` in constructors and destructors is allowed. This is allowed only if anachronisms are enabled and the assignment to this configuration parameter is enabled.
- A bound function pointer (a pointer to a member function for a given object) can be cast to a pointer to a function.

- A nested class name may be used as an unnested class name provided no other class of that name has been declared. The anachronism is not applied to template classes.
- A reference to a non-`const` type may be initialized from a value of a different type. A temporary is created; it is initialized from the (converted) initial value, and the reference is set to the temporary.
- A reference to a non-`const` class type may be initialized from an `rvalue` of the `class` type or a derived class thereof. No (additional) temporary is used.
- A function with old-style parameter declarations is allowed and may participate in function overloading as though it were prototyped. Default argument promotion is not applied to parameter types of such functions when the check for compatibility is done, so that the following statements declare the overload of two functions named `f`:

```
int f(int);
int f(x) char x; { return x; }
```

-eh

The `-eh` (enable exception handling) switch directs the compiler to allow C++ code that contains catch statements and throw expressions and other features associated with ANSI/ISO standard C++ exceptions. When this switch is enabled, the compiler defines the macro `__EXCEPTIONS` to be 1.

The `-eh` switch also causes the compiler to define `__ADI_LIBEH__` during the linking stage so that appropriate sections can be activated in the `.LDF` file, and the program can be linked with a library built with exceptions enabled.

Object files created with exceptions enabled may be linked with objects created without exceptions. However, exceptions can only be thrown from and caught, and cleanup code executed, in modules compiled with `-eh`.

-no-anach

The `-no-anach` (disable C++ anachronisms) switch directs the compiler to disallow some old C++ language features that are prohibited by the C++ standard. See the `-anach` switch (on page 1-67) for a full description of these features.

-no-demangle

The `-no-demangle` (disable demangler) switch directs the compiler to prevent the driver from filtering linker errors through the demangler. The demangler's primary role is to convert the encoded name of a function into a more understandable version of the name.

-no-eh

The `-no-eh` (disable exception handling) directs the compiler to disallow ANSI/ISO C++ exception handling. This is the default mode. See the `-eh` switch (on page 1-68) for more information.

-no-implicit-inclusion

The `-no-implicit-inclusion` switch prevents implicit inclusion of source files as a method of finding definitions of template entities to be instantiated.

-no-rtti

The `-no-rtti` (disable run-time type identification) switch directs the compiler to disallow support for `dynamic_cast` and other features of ANSI/ISO C++ run-time type identification. This is the default mode. Use `-rtti` to enable this feature.

-rtti

The `-rtti` (enable run-time type identification) switch directs the compiler to accept programs containing `dynamic_cast` expressions and other features of ANSI/ISO C++ run-time type identification. The switch also causes the compiler to define the macro `__RTTI` to 1. See also the `-no-rtti` switch.

Environment Variables Used by the Compiler

The compiler refers to a number of environment variables during its operation, as listed below. The majority of the environment variables identify *path names* to directories. You should be aware that placing network paths into these environment variables may adversely affect the time required to compile applications.

- **PATH**
  This is your System search path, used to locate Windows applications when you run them. Windows uses this environment variable to locate the compiler when you execute it from the command line.
- **TMP**
  This directory is used by the compiler for temporary files, when building applications. For example, if you compile a C file to an object file, the compiler first compiles the C file to an assembly file which can be assembled to create the object file. The compiler usually creates a temporary directory within the TMP directory into which to put such files. However, if the `-save-temps` switch is specified, the compiler creates temporary files in the current directory instead. This directory should exist and be writable. If this directory does not exist, the compiler issues a warning.
- **TEMP**
  This environment variable is also used by the compiler when looking for temporary files, but only if TMP was examined and was not set or the directory that TMP specified did not exist.

- **ADI_DSP**
  The compiler locates other tools in the tool-chain through the VisualDSP++ installation directory, or through the `-path-install` switch. If neither is successful, the compiler looks in `ADI_DSP` for other tools.
- **CCBLKFN_OPTIONS**
  If this environment variable is set, and `CCBLKFN_IGNORE_ENV` is not set, this environment variable is interpreted as a list of additional switches to be appended to the command line. Multiple switches are separated by spaces or new lines. A vertical-bar (|) character may be used to indicate that any switches following it will be processed after all other command-line switches.
- **CCBLKFN_IGNORE_ENV**
  If this environment variable is set, `CCBLKFN_OPTIONS` is ignored.

Optimization Control

The general aim of compiler optimization is to generate correct code that executes fast and is small in size. Not all optimizations are suitable for every application or could be used all the time. Therefore, the compiler optimizer has a number of configurations, or optimization levels, which can be applied when needed. Each of these levels are enabled by one or more compiler switches (and VisualDSP++ project options) or pragmas.

Refer to Chapter 2, "Achieving Optimal Performance from C/C++ Source Code" for information on how to obtain maximal code performance from the compiler.

The following list identifies several optimization levels. The levels are notionally ordered with least optimization listed first and most optimization listed last. The descriptions for each level outline the optimizations performed by the compiler and identify any switches or pragmas required, or that have direct influence on the optimization levels performed.

- **Debug**
  The compiler produces debug information to ensure that the object code matches the appropriate source code line. See "-g" on page 1-33 for more information.

- **Default**
  The compiler does not perform any optimization by default when none of the compiler optimization switches are used (or enabled in VisualDSP++ project options). Default optimization level can be enabled using the `optimize_off` pragma (on page 1-159).

- **Procedural Optimizations**
  The compiler performs advanced, aggressive optimization on each procedure in the file being compiled. The optimizations can be directed to favor optimizations for speed (`-O1` or `O`) or space (`-Os`) or a factor between speed and space (`-Ov`). If debugging is also requested, the optimization is given priority so the debugging functionality may be limited. See "-O[0|1]" on page 1-45, "-Os" on page 1-47, "-Ov num" on page 1-47, and "-Og" on page 1-46.

  Procedural optimizations for speed and space (`-O` and `-Os`) can be enabled in C/C++ source using the pragma `optimize_{for_speed|for_space}` (see on page 1-159 for more information on optimization pragmas). The `-Ofp` compiler switch directs the compiler to offset the frame pointer if doing so allows more 16-bit instructions to be used. Offsetting the frame pointer means the function does not conform to the Application Binary Interface (ABI), but allows the compiler to produce smaller code, which, in turn, allows for more multi-issue instructions. Since the ABI is affected, the debugger would be unable to interpret the resulting frame structure; therefore, this option is not allowed in conjunction with `-g`. See "-Ofp" on page 1-46 for more information.

- **Profile-Guided Optimizations (PGO)**
  The compiler performs advanced aggressive optimizations using profiler statistics generated from running the application using representative training data. PGO can be used in conjunction with IPA and Automatic Inlining. See "-pguide" on page 1-50 for more information.

  The most common scenario in collecting PGO data is to setup one or more simple File-to-Device streams where the File is a standard ASCII stream input file and the Device is any stream device supported by the simulator target, such as memory and peripherals. The PGO process can be broken down into the execution of one or more "Data Sets" where a Data Set is the association of zero or more input streams with one and only one `.pgo` output file. The user can create, edit and delete the Data Sets through the IDDE and then "run" the Data Sets with the click of one button to produce an optimized application. The PGO operation is handled via a new PGO submenu added to the top-level **Tools** menu: **Tools -> PGO -> Manage Data Sets**.

  For more information on PGO, see the *VisualDSP++ 4.0 Getting Started Guide* and the online Help.

- **Automatic Inlining**
  The compiler automatically inlines C/C++ functions which are not necessarily declared as inline in the source code. It does this when it has determined that doing so reduces execution time. How aggressively the compiler performs automatic inlining is controlled using the `-Ov` switch. Automatic inlining is enabled using the `-Oa` switch and additionally enables Procedural Optimizations (`-O`). See "-Oa" on page 1-46, "-Ov num" on page 1-47, and "-O[0|1]" on page 1-45 for more information.

When remarks are enabled, the compiler produces a remark to indicate each function that is inlined.

- **Interprocedural optimizations**
  The compiler performs advanced, aggressive optimization over the whole program, in addition to the per-file optimizations in procedural optimization. The *interprocedural analysis* (IPA) is enabled using the `-ipa` switch and additionally enables procedural optimizations (`-O`). See "-ipa" on page 1-36, "-O[0|1]" on page 1-45, and "Interprocedural Analysis" for more information.

The compiler optimizer attempts to vectorize loops when it is safe to do so. When IPA is used it can identify additional safe candidates for vectorization which might not be classified as safe at a Procedural Optimization level. Additionally, there may be other loops that are known to be safe candidates for vectorization which can be identified to the compiler with use of various pragmas (see "Loop Optimization Pragmas" on page 1-156).

Using the various compiler optimization levels is an excellent way of improving application performance. However consideration should be given to how applications are written so that compiler optimizations are given the best opportunity to be productive. These issues are the topic of Chapter 2, "Achieving Optimal Performance from C/C++ Source Code".

Interprocedural Analysis

The compiler has an optimization capability called *Interprocedural Analysis* (IPA) that allows the compiler to optimize across translation units instead of within individual translation units. This capability allows the compiler to see all of the source files used in a final link at compilation time and to use that information while optimizing.

Interprocedural analysis is enabled by selecting the **Interprocedural analysis** option on the **Compiler** tab (accessed via the VisualDSP++ **Project Options** dialog box), or by specifying the `-ipa` command-line switch (on page 1-36).

The `-ipa` switch automatically enables the `-O` switch to turn on optimization. However, all object files that are supplied in the final link must have been compiled with the `-ipa` switch; otherwise, undefined behavior may result.

When interprocedural analysis is enabled, the pragmas controlling optimization are disabled (see "General Optimization Pragmas" on page 1-159).

Use of the `-ipa` switch generates additional files along with the object file produced by the compiler. These files have `.ipa` and `.opa` filename extensions and should not be deleted manually unless the associated object file is also deleted.

All of the `-ipa` optimizations are invoked after the initial link, when a special program called the prelinker reinvokes the compiler to perform the new optimizations.

Because a file may be recompiled by the prelinker, do not use the `-S` option to see the final optimized assembler file when `-ipa` is enabled. Instead, use the `-save-temps` switch, so that the full compile/link cycle can be performed first.

Because IPA operates only during the final link, the `-ipa` switch has no benefit when compiling the source files to object format for inclusion in a library. Although IPA generates usage information for potential additional optimizations at the final link stage, neither the usage information nor the module's source file are available when the linker includes a module from a library. Therefore, each library module is compiled to the normal `-O` optimization level.

The prelinker inspects object modules included from libraries and other object files which were not compiled with the `-ipa` switch. This allows to see whether there are hidden references to the functions and variables defined in those objects, which were compiled with the `-ipa` switch, and optimizes those variables and functions accordingly.

Interaction with Libraries

When IPA is enabled, the compiler examines all of the source files to build up usage information about all of the function and data items. It then uses that information to make additional optimizations across all of the source files. One of these optimizations is to remove functions that are never called. This optimization can significantly reduce the overall size of the final executable.

Because IPA operates only during the final link, the `-ipa` switch has no benefit when initially compiling source files to object format for inclusion in a library. Although IPA generates usage information for potential additional optimizations at the final link stage, as normal, neither the usage information nor the module's source file are available when the linker includes a module from a library. Each library module has been compiled to the normal `-O` optimization level, but the prelinker cannot access the previously-generated additional usage information for an object in a library. Therefore, IPA cannot exploit the additional information associated with a library module.

If a library module has to make calls to a function in a user module in the program, IPA must be told that this call may occur. The reason IPA needs to know about the call is because IPA examines all the visible calls to the function and determines how best to optimize it based on that information. However, it cannot "see" the calls to the function from the library because the library code has no associated usage information to show that it uses the function.

There is a pragma, `retain_name`, that tells IPA that there are calls that it cannot see, as shown in the following example.

```
int delete_me(int x) {
    return x-2;
}

#pragma retain_name("keep_me")
int keep_me(int y) {
    return y+2;
}

int main(void) {
    return 0;
}
```

When this program is compiled and linked with the -ipa switch, IPA can see that there are no calls to `delete_me()` in any of the source files (one source file, in this case); therefore, IPA deletes it since it becomes unnecessary. IPA does not delete `keep_me()` because the `retain_name` pragma tells IPA that there are uses of the function not visible to IPA. No pragma is necessary for `main()` because IPA knows this is the entry-point to the program.

IPA assumes that it can see all calls to a function and makes use of its knowledge of the parameters being passed to a function to effectively tailor the code generated for a function. If there are calls on a function from an object module in a library, then IPA will not have access to the information for that invocation of the function; this may cause it to incorrectly optimize the generated code.

C/C++ Compiler Language Extensions

The compiler supports extensions to the ANSI/ISO standard for the C and C++ languages. These extensions add support for DSP hardware and permit some C++ programming features when compiling in C mode. The extensions are also available when compiling in C++ mode.

This section provides an overview of the extensions, brief descriptions, and pointers to more information on each extension.

This section contains:

- "Inline Function Support Keyword (inline)" on page 1-81
- "Inline Assembly Language Support Keyword (asm)" on page 1-82
- "Bank Type Qualifiers" on page 1-96
- "Placement Support Keyword (section)" on page 1-97
- "Placement of Compiler-Generated Data" on page 1-98
- "Boolean Type Support Keywords (bool, true, false)" on page 1-98
- "Pointer Class Support Keyword (restrict)" on page 1-99
- "Non-Constant Aggregate Initializer Support" on page 1-100
- "Indexed Initializer Support" on page 1-100
- "Variable-Length Arrays" on page 1-102
- "C++ Style Comments" on page 1-102
- "Compiler Built-In Functions" on page 1-103
- "Pragmas" on page 1-148
- "GCC Compatibility Extensions" on page 1-187
- "Preprocessor-Generated Warnings" on page 1-195

The additional keywords that are part of the C/C++ extensions do not conflict with ANSI C/C++ keywords. The formal definitions of these extension keywords are prefixed with a leading double underscore (__). Unless the `-no-extra-keywords` command-line switch is used, the compiler defines the shorter form of the keyword extension that omits the leading underscores. For more information, see brief descriptions of each switch beginning on page 1-25.

This section describes the shorter forms of the keyword extensions. In most cases, you can use either form in your code. For example, all references to the `inline` keyword in this text appear without the leading double underscores, but you can interchange `inline` and `__inline` in your code.

You might exclusively use the longer form (such as `__inline`) if porting a program that uses the extra Analog Devices keywords as identifiers. For example, if a program declares local variables, such as `asm` or `inline`, use the `-no-extra-keywords` switch. If you need to declare a function as inline, you can use `__inline`.

Table 1-7 and Table 1-8 provide descriptions of each extension and direct you to sections that describe each extension in more detail.

Table 1-7. Keyword Extensions

Keyword Extensions	Description
`inline`	Directs the compiler to integrate the function code into the code of its callers. For more information, see "Inline Function Support Keyword (inline)" on page 1-81.
`asm()`	Places Blackfin core assembly language commands directly in your C/C++ program. For more information, see "Inline Assembly Language Support Keyword (asm)" on page 1-82.
`bank("string")`	Specifies a name which the user assigns to associate declarations that reside in particular memory banks. For more information, see "Bank Type Qualifiers" on page 1-96.

Table 1-7. Keyword Extensions (Cont'd)

Keyword Extensions	Description
`section("string")`	Specifies the section in which an object or function is placed. For more information, see "Placement Support Keyword (section)" on page 1-97.
`bool`, `true`, `false`	Specifies a Boolean type. For more information, see "Boolean Type Support Keywords (bool, true, false)" on page 1-98.
`restrict`	Specifies restricted pointer features. For more information, see "Pointer Class Support Keyword (restrict)" on page 1-99.

Table 1-8. Operational Extensions

Operational Extensions	Description
Non-constant initializers	Lets you use non-constants as elements of aggregate initializers for automatic variables. For more information, see "Non-Constant Aggregate Initializer Support" on page 1-100.
Indexed initializers	Lets you specify elements of an aggregate initializer in arbitrary order. For more information, see "Indexed Initializer Support" on page 1-100.
Variable-length arrays	Lets you create local arrays with a variable size. For more information, see "Variable-Length Arrays" on page 1-102.
Preprocessor-generated warnings	Lets you generate warning messages from the preprocessor. For more information, see "Preprocessor-Generated Warnings" on page 1-195.
C++ style comments	Allows for "//" C++ style comments in C programs. For more information, see "C++ Style Comments" on page 1-102.

Inline Function Support Keyword (inline)

The `inline` keyword directs the compiler to integrate the code for the function you declare as `inline` into the code of its callers. Inline function support and the `inline` keyword is a standard feature of C++; the `ccblkfn` compiler provides this keyword as a C extension.

This keyword eliminates the function call overhead and increases the speed of your program's execution. Argument values that are constant and that have known values may permit simplifications at compile time so that not all of the inline function's code needs to be included. The following example shows a function definition that uses the `inline` keyword.

```
inline int max3 (int a, int b, int c) {
   return max (a, max(b, c));
}
```

A function declared `inline` must be defined (its body must be included) in every file in which the function is used. This is normally done by placing the `inline` definition in a header file. Usually it is also declared static.

In some cases, the compiler does not output object code for the function; for example, when the address is not needed for an `inline` function which is called from within the defining program. However, recursive calls and functions whose addresses are explicitly referred to by the program are compiled to assembly code.

The compiler only inlines functions, even those declared using the `inline` keyword, when optimizations are enabled (using the `-O` switches, as described on page 1-45).

Inline Assembly Language Support Keyword (asm)

The compiler's `asm()` construct is used to code Blackfin assembly language instructions within a C/C++ function and to pass declarations and directives through to the assembler. Use the `asm()` construct to express assembly language statements that cannot be expressed easily or efficiently with C or C++ constructs.

Using `asm()`, you can code complete assembly language instructions and specify the operands of the instruction using C expressions. When specifying operands with a C/C++ expression, you do not need to know which registers or memory locations contain C/C++ variables.

The compiler *does not analyze* code defined with the `asm()` construct—it passes this code directly to the assembler. The compiler *does* perform substitutions for operands of the formats `%0` through `%9`. However, it passes *everything else* to the assembler without reading or analyzing it.

The `asm()` constructs are executable statements, and as such, may not appear before declarations within C/C++ functions.

The `asm()` constructs may also be used at global scope, outside function declarations. Such `asm()` constructs are used to pass declarations and directives directly to the assembler. They are not executable constructs, and may not have any inputs or outputs, or affect any registers.

When optimizing, the compiler sometimes changes the order in which generated functions appear in the output assembly file. However, if global-scope `asm` constructs are placed between two function definitions, the compiler ensures that the function order is retained in the generated assembly file. Consequently, function inlining may be inhibited.

A simplified `asm()` construct without operands takes the form of

```
asm(" nop; ");
```

The complete assembly language instruction, enclosed in double quotes, is the argument to `asm()`. Using `asm()` constructs with operands requires some additional syntax.

The compiler generates a label before and after inline assembly instructions when generating debug code (the `-g` switch on page 1-33). These labels are used to generate the debug line information used by the debugger. If the inline assembler inserts conditionally assembled code, an undefined symbol error is likely to occur at link time. For example, the following code could cause undefined symbols if `MACRO` is undefined:

```
asm("#ifdef MACRO");
asm(" // assembly statements");
asm("#endif");
```

If the inline assembler changes the current section and thereby causes the compiler labels to be placed in another section, such as a data section (instead of the default code section), then the debug line information will be incorrect for these lines.

The construct syntax is described in:

- "Assembly Construct Template" on page 1-84
- "Assembly Construct Operand Description" on page 1-87
- "Assembly Constructs With Multiple Instructions" on page 1-93
- "Assembly Construct Reordering and Optimization" on page 1-94
- "Assembly Constructs With Input and Output Operands" on page 1-94
- "Assembly Constructs and Flow Control" on page 1-95
- "Guildelines on the Use of asm() Statements" on page 1-96

Assembly Construct Template

Use `asm()` constructs to specify the operands of assembly instruction using C or C++ expressions. You do not need to know which registers or memory locations contain C or C++ variables.

asm() Constructs Syntax

Use the following general syntax for your `asm()` constructs.

```
asm(
   template
   [:[constraint(output operand)[,constraint(output operand)…]]
     [:[constraint(input operand)[,constraint(input operand)…]]
            [:clobber]]]
);
```

The syntax elements are defined as:

template

The template is a string containing the assembly instruction(s) with `%number` indicating where the compiler should substitute the operands. Operands are numbered in order of occurrence from left to right, starting at 0. Separate multiple instructions with a semicolon; then enclose the entire string within double quotes.
For more information on templates containing multiple instructions, see "Assembly Constructs With Multiple Instructions" on page 1-93.

constraint

The constraint is a string that directs the compiler to use certain groups of registers for the input and output operands. Enclose the constraint string within double quotes. For more information on operand constraints, see "Assembly Construct Operand Description" on page 1-87.

output operand

The output operands are the names of a C/C++ variables that receives output from a corresponding operand in the assembly instructions.

input operand

The input operand is a C/C++ expression that provides an input to a corresponding operand in the assembly instruction.

clobber

The clobber notifies the compiler that a list of registers is overwritten by the assembly instructions. Use lowercase characters to name clobbered registers. Enclose each name within double quotes, and separate each quoted register name with a comma. The input and output operands are guaranteed not to use any of the clobbered registers, so you can read and write the clobbered registers as often as you like. See Table 1-10 on page 1-92.

asm() Construct Syntax Rules

These rules apply to assembly construct template syntax.

- The template is the only mandatory argument to `asm()`. All other arguments are optional.
- An operand constraint string followed by a C/C++ expression in parentheses describes each operand. For output operands, it must be possible to assign to the expression; that is, the expression must be legal on the left side of an assignment statement.
- A colon separates:
  - The template from the first output operand
  - The last output operand from the first input operand
  - The last input operand from the clobbered registers

- A space must be added between adjacent colon field delimiters in order to avoid a clash with the C++ "::" reserved global resolution operator.
- A comma separates operands and registers within arguments.
- The number of operands in arguments must match the number of operands in your template.
- The maximum permissible number of operands is ten (`%0`, `%1`, `%2`, `%3`, `%4`, `%5`, `%6`, `%7`, `%8`, and `%9`).

The compiler cannot check whether the operands have data types that are reasonable for the instruction being executed. The compiler does not parse the assembler instruction template, does not interpret the template, and does not verify whether the template contains valid input for the assembler.

asm() Construct Template Example

The following example shows how to apply the `asm()` construct template to the Blackfin assembly language assignment instruction.

```
{
int result, x;
...
asm (
   "%0=%1;" :
   "=d" (result) :
   "d" (x)
   );
}
```

In the previous example, note that:

- The template is `"%0=%1;"`. The `%0` is replaced with operand zero `(result)`; the first operand, `%1`, is replaced with operand one `(x)`.

- The output operand is the C/C++ variable `result`. The letter `d` is the operand constraint for the variable. This constrains the output to a data register `R{0-7}`. The compiler generates code to copy the output from the `r` register to the variable `result`, if necessary. The `=` in `=d` indicates that the operand is an output.
- The input operand is the C/C++ variable `x`. The letter `d` in the operand constraint position for this variable constrains `x` to a data register `R{0-7}`. If `x` is stored in a different kind of register or in memory, the compiler generates code to copy the value into an `r` register before the `asm()` construct uses it.

Assembly Construct Operand Description

The second and third arguments to the `asm()` construct describe the operands in the assembly language template. Several pieces of information must be conveyed for the compiler to know how to assign registers to operands. This information is conveyed with an operand constraint. The compiler needs to know what kind of registers the assembly instructions can operate on, so it can allocate the correct register type.

You convey this information with a letter in the operand constraint string that describes the class of allowable registers.

Table 1-9 on page 1-91 describes the correspondence between constraint letters and register classes.

The use of any letter not listed in Table 1-9 results in unspecified behavior. The compiler does not check the validity of the code by using the constraint letter.

To assign registers to the operands, the compiler must also be informed of which operands in an assembly language instruction are inputs, which are outputs, and which outputs may not overlap inputs. The compiler is told this in three ways.

- The output operand list appears as the first argument after the assembly language template. The list is separated from the assembly language template with a colon. The input operands are separated from the output operands with a colon and they always follow the output operands.
- The operand constraints describe which registers are modified by an assembly language instruction. The "=" in `=constraint` indicates that the operand is an output; all output operand constraints must use =. Operands that are input-outputs must use "+" (see below).
- The compiler may allocate an output operand in the same register as an unrelated input operand, unless the output or input operand has the `&=` constraint modifier. This situation can occur because the compiler assumes the inputs are consumed before the outputs are produced.

  This assumption may be false if the assembler code actually consists of more than one instruction. In such a case, use `&=` for each output operand that must not overlap an input or supply an `&` for the input operand.

Operand constraints indicate what kind of operand they describe by means of preceding symbols. The possible preceding symbols are: no symbol, =, +, &, ?, and #.

- (no symbol)

  The operand is an input. It must appear as part of the third argument to the `asm()` construct. The allocated register is loaded with the value of the C/C++ expression before the `asm()` template is executed. Its C/C++ expression is not modified by the `asm()` construct, and its value may be a constant or literal.
  Example: `d`

- **= symbol**

  The operand is an output. It must appear as part of the second argument to the `asm()` construct. Once the `asm()` template has been executed, the value in the allocated register is stored into the location indicated by its C/C++ expression; therefore, the expression must be one that would be valid as the left-hand side of an assignment.
  Example: `=d`

- **+ symbol**

  The operand is both an input and an output. It must appear as part of the second argument to the `asm()` construct. The allocated register is loaded with the C/C++ expression value, the `asm()` template is executed, and then the allocated register's new value is stored back into the C/C++ expression. Therefore, as with pure outputs, the C/C++ expression must be one that is valid on the left-hand side of an assignment.
  Example: `+d`

- **? symbol**

  The operand is temporary. It must appear as part of the third argument to the `asm()` construct. A register is allocated as working space for the duration of the `asm()` template execution. The register's initial value is undefined, and the register's final value is discarded. The corresponding C/C++ expression is not loaded into the register, but must be present. This expression is normally specified using a literal zero.
  Example: `?d`

- **& symbol**

  This operand constraint may be applied to inputs and outputs. It indicates that the register allocated to the input (or output) may not be one of the registers that are allocated to the outputs (or inputs). This operand constraint is used when one or more output registers are set while one or more inputs are yet to be referenced. (This situation sometimes occurs if the `asm()` template contains more than one instruction.)
  Example: `&d`

- **# symbol**

  The operand is an input, but the register's value is clobbered by the `asm()` template execution. The compiler may make no assumptions about the register's final value. An input operand with this constraint will not be allocated the same register as any other input or output operand of the `asm()`. The operand must appear as part of the second argument to the `asm()` construct.
  Example: `#d`

Table 1-9 lists the registers that may be allocated for each register constraint letter. The use of any letter not listed in the "Constraint" column of this table results in unspecified behavior. The compiler does not check the validity of the code by using the constraint letter. Table 1-10 lists the registers that may be named as part of the clobber list.

It is also possible to claim registers directly, instead of requesting a register from a certain class using the constraint letters. You can claim the registers directly by simply naming the register in the location where the class letter would be. The register names are the same as those used to specify the clobber list; see Table 1-10.

For example,

```
asm("%0 += %1 * %2;"
:"+a0"(sum)     /* output */
:"H"(x),"H"(y)  /* input  */
);
```

would load `sum` into `A0`, and load `x` and `y` into two `DREG` halves, execute the operation, and then store the new total from `A0` back into `sum`.

Naming the registers in this way allows the `asm()` construct to specify several registers that must be related, such as the `DAG` registers for a circular buffer. This also allows the use of registers not covered by the register classes accepted by the `asm()` construct. The clobber string can be any of the registers recognized by the compiler.

Table 1-9. asm() Operand Constraints

Constraint	Register Type	Registers
a	General addressing registers	P0 – P5
p	General addressing registers	P0 – P5
i	DAG addressing registers	I0 – I3
b	DAG addressing registers	I0 – I3
d	General data registers	R0 – R7
r	General data registers	R0 – R7
D	General data registers	R0 – R7
A	Accumulator registers	A0, A1
e	Accumulator registers	A0, A1
f	Modifier registers	M0 – M3
E	Even general data registers	R0,R2,R4,R6
O	Odd general data registers	R1,R3,R5,R7
h	High halves of the general data registers	R0.H,R1.H...R7.H

Table 1-9. asm() Operand Constraints (Cont'd)

Constraint	Register Type	Registers
l	Low halves of the general data registers	R0.L,R1.L...R7.L
H	Low or high halves of the general data registers	R0.L,R1.L...R7.L
L	Loop counter registers	LC0,LC1
constraint	Indicates the constraint is an input operand	
=constraint	Indicates the constraint is applied to an output operand	
&constraint	Indicates the constraint is applied to an input operand that may not be overlapped with an output operand	
=&constraint	Indicates the constraint is applied to an output operand that may not overlap an input operand	
?constraint	Indicates the constraint is temporary	
+constraint	Indicates the constraint is both an input and output operand	
#constraint	Indicates the constraint is an input operand whose value will be changed	

Table 1-10. Register Names for asm() Constructs

Clobber String	Meaning
"r0", "r1", "r2", "r3", "r4", "r5", "r6", "r7",	General data register
"p0", "p1", "p2", "p3", "p4", "p5",	General addressing register
"i0","i1", "i2", "i3",	DAG addressing register
"m0", "m1", "m2", "m3",	Modifier register
"b0", "b1", "b2", "b3",	Base register
"l0", "l1", "l2", "l3",	Length register
"astat",	ALU status register
"seqstat",	Sequencer status register
"rets",	Subroutine address register
"cc",	Condition code register

Table 1-10. Register Names for asm() Constructs (Cont'd)

Clobber String	Meaning
"a0", "a1",	Accumulator result register
"lc0", "lc1",	Loop counter register
"memory"	Unspecified memory location(s)

Assembly Constructs With Multiple Instructions

There can be many assembly instructions in one template. Your template string may extend over multiple lines. It is not necessary to terminate each line with a backslash (\).

This is an example of multiple instructions in a template:

```
/* (pseudo code) r7 = x; r6 = y; result = x + y; */
   asm ("r7=%1;
   r6=%2;
   %0=r6+r7;"
   : "=d" (result)              /* output    */
   : "d" (x), "d" (y)           /* input     */
   : "r7", "r6");               /* clobbers  */
```

Do not attempt to produce multiple-instruction `asm` constructs via a sequence of single-instruction `asm` constructs, as the compiler is not guaranteed to maintain the ordering. For example, the following should be avoided:

```
   /* BAD EXAMPLE: Do not use sequences of single-instruction
   ** asms. Use a single multiple-instruction asm instead.  */

   asm("r7=%0;" : : "d" (x) : "r7");
   asm("r6=%0;" : : "d" (y) : "r6");
   asm("%0=r6+r7;" : "=d" (result));
```

Assembly Construct Reordering and Optimization

For the purpose of optimization, the compiler assumes that the side effects of an `asm()` construct are limited to changes in the output operands or the items specified using the clobber specifiers. This does not mean that you cannot use instructions with side effects, but be careful to notify the compiler that you are using them by using the clobber specifiers (see Table 1-10).

The compiler may eliminate supplied assembly instructions if the output operands are not used, move them out of loops, or replace two with one if they constitute a common subexpression. Also, if the instruction has a side effect on a variable that otherwise appears not to change, the old value of the variable may be reused later if it happens to be found in a register.

Use the keyword `volatile` to prevent an `asm()` instruction from being moved, combined, or deleted. For example,

```
#define set_priority(x) \
asm volatile ("STI %0;":   /* no outs */ : "d" (x))
```

A sequence of `asm volatile()` constructs is not guaranteed to be completely consecutive; it may be moved across jump instructions or in other ways that are not significant to the compiler. To force the compiler to keep the output consecutive, use one `asm volatile()` construct only, or use the output of the `asm()` construct in a C/C++ statement.

Assembly Constructs With Input and Output Operands

The output operands must be write only; the compiler assumes that the values in these operands do not need to be preserved. When the assembler instruction has an operand that is read from and written to, you must logically split its function into two separate operands: one input operand and one write-only output operand. The connection between the two operands is expressed by constraints in the same location when the instruction executes.

When a register's value is both an input and an output, and the final value is to be stored to the same location from which the original value was loaded, the register can be marked as an input-output, using the "+" constraint symbol, as described earlier.

If the final value is to be saved into a different location, then both an input and an output must be specified, and the input must be tied to the output by using its position number as the constraint.

For example,

```
asm("%0 += 4;"
   :"=p" (newptr)      // an output, given a preg,
                       // stored into newptr.
   :"0" (oldptr));     // an input, given same reg as %0,
                       // initialized from oldptr
```

Assembly Constructs and Flow Control

Do not place flow control operations within an `asm()` construct that "leaves" the `asm()` construct functions, such as calling a procedure or performing a jump, to another piece of code that is not within the `asm()` construct itself. Such operations are invisible to the compiler, may result in multiple-defined symbols, and may violate assumptions made by the compiler.

For example, the compiler is careful to adhere to the calling conventions for preserved registers when making a procedure call. If an `asm()` construct calls a procedure, the `asm()` construct must also ensure that all conventions are obeyed, or the called procedure may corrupt the state used by the function containing the `asm()` construct.

It is also inadvisable to use labels in `asm()` statements, especially when function inlining is enabled. If a function containing such `asm` statements is inlined more than once in a file, there will be multiple definitions of the label, resulting in an assembler error. If possible, use PC-relative jumps in `asm` statements.

Guildelines on the Use of asm() Statements

There are certain operations that are performed more efficiently using other compiler features, and result in source code that is clearer and easier to read.

Accessing System Registers
System registers are accessed most efficiently using the functions in `sysreg.h` instead of using `asm()` statements.

Accessing Memory-Mapped Registers (MMRs)
MMRs can be accessed using the macros in the `Cdef*.h` files (for example, `cdefBF531.h`) that are supplied with VisualDSP++.

Bank Type Qualifiers

Bank qualifiers can be attached to data declarations to indicate that the data resides in particular memory banks. For example,

```
int bank("blue") *ptr1;
int bank("green") *ptr2;
```

The bank qualifier assists the optimizer because the compiler assumes that if two data items are in different banks, they can be accessed together without conflict.

The bank name string literals have no significance, except to differentiate between banks. There is no interpretation of the names attached to banks, which can be any arbitrary string. There is a current implementation limit of ten different banks.

For any given function, three banks are automatically defined. These are:

- The default bank for global data.
  The "static" or "extern" data that is not explicitly placed into another bank is assumed to be within this bank. Normally, this bank is called "`__data`"", although a different bank can be selected with `#pragma data_bank(bankname)`.

- The default bank for local data.
  Local variables of "auto" storage class that are not explicitly placed into another bank are assumed to be within this bank. Normally, this bank is called "__stack", although a different bank can be selected with #pragma stack_bank(bankname).
- The default bank for the function's instructions.
  The function itself is placed into this bank. Normally, it is called "__code", although a different bank can be selected with #pragma code_bank(bankname).

Each memory bank can have different performance characteristics. For more information on memory bank attributes, see "Memory Bank Pragmas" on page 1-181.

Placement Support Keyword (section)

The section() keyword directs the compiler to place an object or function in an assembly .SECTION of the compiler's intermediate output file. You name the assembly .SECTION with the string literal parameter of the section() keyword. If you do not specify a section() keyword for an object or function declaration, the compiler uses a default section. The .LDF file supplied to the linker must also be updated to support the additional named section. For information on the default sections, see "Using Memory Sections" on page 1-243.

Applying section() is meaningful only when the data item is something that the compiler can place in the named section. Apply section() only to top-level, named objects that have static duration; for example, objects that are explicitly static, or are given as external-object definitions. The example below shows the declaration of a static variable that is placed in the section called bingo.

```
static section("bingo") int x;
```

Placement of Compiler-Generated Data

The `-section` switch (on page 1-54) can be used to control the placement of compiler-generated data into non-default sections.

When using certain features of C/C++, the compiler may be required to produce internal data structures. The `section` switch allows you to override the default location where the data would be placed. For example,

```
ccblkfn -section vtbl=vtbl_data test.cpp -c++
```

would instruct the compiler to place all the C++ virtual function look-up tables into the section `vtbl_data`, rather than the default `constdata` section. It is the user's responsibility to ensure that appropriately named sections exist in the `.LDF` file.

The following section identifiers are currently supported:

`vtbl` C++ virtual function lookup tables

Boolean Type Support Keywords (bool, true, false)

The `bool`, `true`, and `false` keywords are extensions that support the C++ boolean type in C mode. The `bool` keyword is a unique signed integral type, just as the `wchar_t` is a unique unsigned type. There are two built-in constants of this type: `true` and `false`. When converting a numeric or pointer value to `bool`, a zero value becomes `false`, and a nonzero value becomes `true`. A `bool` value may be converted to `int` by promotion, taking `true` to one and `false` to zero. A numeric or pointer value is automatically converted to `bool` when needed.

These keyword extensions behave as if the declaration that follows had appeared at the beginning of the file, except that assigning a nonzero integer to a `bool` type always causes it to take on the value `true`.

```
typedef enum { false, true } bool;
```

Pointer Class Support Keyword (restrict)

The `restrict` keyword is an extension that supports restricted pointer features. The use of `restrict` is limited to the declaration of a pointer. This keyword specifies that the pointer provides exclusive initial access to the pointed object. More simply, the `restrict` keyword is a way to identify that a pointer does not create an alias. Also, two different restricted pointers cannot designate the same object and therefore they are not aliases.

The compiler is free to use the information about restricted pointers and aliasing in order to better optimize C or C++ code that uses pointers. The `restrict` keyword is most useful when applied to function parameters that the compiler otherwise have little information about.

```
void fir(short *in, short *c, short *restrict out, int n)
```

The behavior of a program is undefined if it contains an assignment between two restricted pointers. Exceptions are:

- A function with a restricted pointer parameter may be called with an argument that is a restricted pointer.
- A function may return the value of a restricted pointer that is local to the function, and the return value may then be assigned to another restricted pointer.

If your program uses a restricted pointer in a way that it does not uniquely refer to storage, the behavior of the program is undefined.

Non-Constant Aggregate Initializer Support

The compiler includes extended support for aggregate initializers. The compiler does not require the elements of an aggregate initializer for an automatic variable to be constant expressions. The following example shows an initializer with elements that vary at run time.

```
void initializer (float a, float b)
{
   float the_array[2] = { a-b, a+b };
}
```

All automatic structures can be initialized by arbitrary expressions involving literals, previously declared variables and functions.

Indexed Initializer Support

ANSI/ISO Standard C/C++ requires that the elements of an initializer appear in a fixed order—the same order as the elements in the array or structure being initialized. The `ccblkfn` compiler, by comparison, supports labeling elements for array initializers. This feature lets you specify the array or structure elements in any order by specifying the array indices or structure field names to which they apply.

For an array initializer, the syntax `[INDEX]` appearing before an initializer element value specifies the index initialized by that value. Subsequent initializer elements are then applied to the sequentially following elements of the array, unless another use of the `[INDEX]` syntax appears. The index values must be constant expressions, even if the array being initialized is automatic.

The following example shows equivalent array initializers—the first in standard C and C++ and the next using the extension. Note that the [`index`] precedes the value being assigned to that element.

```
/* Example 1 C Array Initializer */
/* Standard C array initializer  */
```

```
int a[6] = { 0, 0, 15, 0, 29, 0 };

/* equivalent ccblkfn C array initializer */

int a[6] = { [4] 29, [2] 15 };
```

You can combine this technique of naming elements with Standard C/C++ initialization of successive elements. The Standard C/C++ and compiler instructions below are equivalent. Note that any unlabeled initial value is assigned to the next consecutive element of the structure or array.

```
/* Example 2 Standard C & ccblkfn  /C++ C Array Initializer */
/* Standard C array initializer */

int a[6] = { 0, v1, v2, 0, v4, 0 };

/* equivalent ccblkfn C array initializer that uses indexed
   elements */

 int a[6] = { [1] v1, v2, [4] v4 };
```

The following example shows how to label the array initializer elements when the indices are characters or `enum` type.

```
/* Example 3 C Array Initializer With enum Type Indices */
/* ccblkfn C array initializer */

int whitespace[256] =
{
[' '] 1, ['\t'] 1, ['\v'] 1, ['\f'] 1, ['\n'] 1, ['\r'] 1
};
```

In a structure initializer, specify the name of the field to initialize with `fieldname:` before the element value. The Standard C/C++ and compiler's C struct initializers in the example below are equivalent.

```
/* Example 4 Standard C & ccblkfn C struct Initializer */
/* Standard C struct Initializer */

struct point {int x, y;};
struct point p = {xvalue, yvalue};
```

```
/* Equivalent ccblkfn C struct Initializer With Labeled
   Elements */

struct point {int x, y;};
struct point p = {y: yvalue, x: xvalue};
```

Variable-Length Arrays

The `ccblkfn` compiler support variable-length automatic arrays in C source (variable length arrays are not suppoted in C++). Standard C requires the size of an array to be known at compile time. The following example shows a function that has an array whose size is determined by the value of a parameter passed onto the function.

```
void var_array (int nelms, int *ival)
{
int temp[nelms];
int i;

for (i=0;i<nelms; i++)
   temp[i] = ival[i]*2;
}
```

Variable-length arrays are only supported as an extension to C and not C++.

C++ Style Comments

The compiler accepts C++ comments, beginning with `//` and ending at the end of the line, as in C programs. This comment representation is essentially compatible with standard C, except for the following case.

```
a = b
//* highly unusual */ c
;
```

which a standard C compiler processes as:

```
a = b/c;
```

and a C++ compiler and `ccblkfn` process as:

```
a = b;
```

Compiler Built-In Functions

The compiler supports intrinsic (built-in) functions that enable efficient use of hardware resources. These functions are:

- "Fractional Value Built-in Functions in C" on page 1-104
- "ETSI Support" on page 1-116
- "Fractional Literal Values in C" on page 1-129
- "Complex Fractional Built-In Functions in C" on page 1-130
- "Complex Operations in C++" on page 1-131
- "Viterbi History and Decoding Functions" on page 1-132
- "Circular Buffer Built-In Functions" on page 1-134
- "Endian-Swapping Intrinsics" on page 1-136
- "System Built-In Functions" on page 1-137
- "Video Operation Built-In Functions" on page 1-140
- "Misaligned Data Built-In Functions" on page 1-147

Knowledge of these functions is built into the `ccblkfn` compiler. Your program uses them via normal function call syntax. The compiler notices the invocation and generates one or more machine instructions, just as it does for normal operators, such as `+` and `*`.

Built-in functions have names which begin with `__builtin_`. Note that identifiers beginning with double underlines (__) are reserved by the C standard, so these names will not conflict with user program identifiers.

The header files also define more readable names for the built-in functions without the `__builtin_` prefix. These additional names are disabled if the `-no-builtin` option is used.

These functions are specific to individual architectures and this section lists the built-in functions supported at this time on Blackfin processors. Various system header files provide definitions and access to the intrinsics as described below.

Fractional Value Built-in Functions in C

The built-in functions provide access to the fractional arithmetic and the parallel 16-bit operations supported by the Blackfin processor instructions. Various C types are defined (Table 1-11) to represent these classes of data.

Table 1-11. Fractional Value C Types

C type	Usage
fract16	Single 16-bit signed fractional value
fract32	Single 32-bit signed fractional value
fract2x16	Double 16-bit signed fractional value

See the section "Using Data Storage Formats" on page 1-261 for more information on how `fract16`, `fract32` and `fract2x16` types are represented. See the instruction set reference of the appropriate Blackfin processor for information on saturation, rounding (biased and unbiased), and truncating.

Because fractional arithmetic uses slightly different instructions to normal arithmetic, you cannot normally use the standard C operators on `fract` data types and get the right result. You should use the built-in functions described here to work with fractional data.

The fract.h header file provides access to the definitions for each of the built-in functions that support fractional values. These functions have names with suffixes _fr1x16 for single fract16, _fr2x16 for dual fract16, and _fr1x32 for single fract32. All the functions in fract.h are marked as inline, so when compiling with the compiler optimizer, the built-in functions are inlined.

All the 16-bit fractional shift built-in functions without "_clip" in the name ignore all but the least significant five bits of the shift magnitude. All the 32-bit fractional shift built-in functions without "_clip" in the name ignore all but the least significant 6 bits of the shift magnitude. The _clip variants of these built-in functions automatically clip the shift magnitude to within a 5- or 6-bit range. Use these built-in functions unless the shift magnitude is known to be within the 5- or 6- bit range to avoid unexpected results.

See "fract16 Built-in Functions" for the description of the built-in functions that work primarily with fract16 data. See "fract32 Built-in Functions" on page 1-108 for the description of the built-in functions that work primarily with fract32 data.

See "fract2x16 Built-in Functions" on page 1-110 for the description of the built-in functions that work primarily with fract2x16 data. Note that when compiling programs that use the single data fract16 operations, the compiler optimizer attempts to automatically detect cases where parallel operations can be performed. In other words, re-coding an algorithm to make explicit use of the fract2x16 built-in functions in place of fract1x16 ones does not always yield a performance benefit.

See "ETSI Built-in Functions" on page 1-115 on how to map the European Telecommunications Standards Institute (ETSI) fract functions onto the compiler built-in functions.

fract16 Built-in Functions

All the built-in functions described here are saturating unless otherwise stated. These builtins operate primarily on the `fract16` type, although one of the multiplies returns a `fract32`.

The following built-in functions are available.

`fract16 add_fr1x16(fract16 f1,fract16 f2)`
Performs 16-bit addition of the two input parameters (`f1+f2`)

`fract16 sub_fr1x16(fract16 f1,fract16 f2)`
Performs 16-bit subtraction of the two input parameters (`f1-f2`)

`fract16 mult_fr1x16(fract16 f1,fract16 f2)`
Performs 16-bit multiplication of the input parameters (`f1*f2`).
The result is truncated to 16 bits.

`fract16 multr_fr1x16(fract16 f1,fract16 f2)`
Performs a 16-bit fractional multiplication (`f1*f2`) of the two input parameters. The result is rounded to 16 bits. Whether the rounding is biased or unbiased depends on what the `RND_MOD` bit in the `ASTAT` register is set to.

`fract32 mult_fr1x32(fract16 f1,fract16 f2)`
Performs a fractional multiplication on two 16-bit fractions, returning the 32-bit result. There will be loss of precision.

`fract16 abs_fr1x16(fract16 f1)`
Returns the 16-bit value that is the absolute value of the input parameter. Where the input is `0x8000`, saturation occurs and `0x7fff` is returned.

`fract16 min_fr1x16(fract16 f1, fract16 f2)`
Returns the minimum of the two input parameters.

`fract16 max_fr1x16(fract16 f1, fract16 f2)`
Returns the maximum of the two input parameters.

`fract16 negate_fr1x16(fract16 f1)`
Returns the 16-bit result of the negation of the input parameter (`-f1`). If the input is `0x8000`, saturation occurs and `0x7fff` is returned.

`fract16 shl_fr1x16(fract16 src, short shft)`
Arithmetically shifts the `src` variable left by `shft` places. The empty bits are zero-filled. If `shft` is negative, the shift is to the right by `abs(shft)` places with sign extension.

`fract16 shl_fr1x16_clip(fract16 src, short shft)`
Arithmetically shifts the `src` variable left by `shft` (clipped to 5 bits) places. The empty bits are zero filled. If `shft` is negative, the shift is to the right by `abs(shft)` places with sign extension.

`fract16 shr_fr1x16(fract16 src, short shft)`
Arithmetically shifts the `src` variable right by `shft` places with sign extension. If `shft` is negative, the shift is to the left by `abs(shft)` places, and the empty bits are zero-filled.

`fract16 shr_fr1x16_clip(fract16 src, short shft)`
Arithmetically shifts the `src` variable right by `shft` (clipped to 5 bits) places with sign extension. If `shft` is negative, the shift is to the left by `abs(shft)` places, and the empty bits are zero-filled.

`fract16 shrl_fr1x16(fract16 src, short shft)`
Logically shifts a `fract16` right by `shft` places. There is no sign extension and no saturation – the empty bits are zero-filled.

`fract16 shrl_fr1x16_clip(fract16 src, short shft)`
Logically shifts a `fract16` right by `shft` (clipped to 5 bits) places. There is no sign extension and no saturation – the empty bits are zero-filled.

`int norm_fr1x16(fract16 f1)`
Returns the number of left shifts required to normalize the input variable so that it is either in the interval `0x4000` to `0x7fff`, or in the interval `0x8000` to `0xc000`. In other words,

```
fract16 x;
shl_fr1x16(x,norm_fr1x16(x));
```

returns a value in the range `0x4000` to `0x7fff`, or in the range `0x8000` to `0xc000`.

fract32 Built-in Functions

All the built-in functions described here are saturating unless otherwise stated. These built-in functions operate primarily on the `fract32` type, although there are a couple of functions that convert from `fract32` to `fract16`.

`fract32 add_fr1x32(fract32 f1,fract32 f2)`
Performs 32-bit addition of the two input parameters (`f1+f2`).

`fract32 sub_fr1x32(fract32 f1,fract32 f2)`
Performs 32-bit subtraction of the two input parameters (`f1-f2`).

`fract32 mult_fr1x32x32(fract32 f1,fract32 f2)`
Performs 32-bit multiplication of the input parameters (`f1*f2`).
The result (which is calculated internally with an accuracy of 40 bits) is rounded (biased rounding) to 32 bits.

`fract32 mult_fr1x32x32NS(fract32 f1, fract32 f2)`
Performs 32-bit non-saturating multiplication of the input parameters (`f1*f2`). This is somewhat faster than `mult_fr1x32x32`. The result (which is calculated internally with an accuracy of 40 bits) is rounded (biased rounding) to 32 bits.

`fract32 abs_fr1x32(fract32 f1)`
Returns the 32-bit value that is the absolute value of the input parameter. Where the input is `0x80000000`, saturation occurs and `0x7fffffff` is returned.

`fract32 min_fr1x32(fract32 f1, fract32 f2)`
Returns the minimum of the two input parameters

`fract32 max_fr1x32(fract32 f1, fract32 f2)`
Returns the maximum of the two input parameters

`fract32 negate_fr1x32(fract32 f1)`
Returns the 32-bit result of the negation of the input parameter (`-f1`). If the input is `0x80000000`, saturation occurs and `0x7fffffff` is returned.

`fract32 shl_fr1x32(fract32 src, short shft)`
Arithmetically shifts the `src` variable left by `shft` places. The empty bits are zero filled. If `shft` is negative, the shift is to the right by `abs(shft)` places with sign extension.

`fract32 shl_fr1x32_clip(fract32 src, short shft)`
Arithmetically shifts the `src` variable left by `shft` (clipped to 6 bits) places. The empty bits are zero filled. If `shft` is negative, the shift is to the right by `abs(shft)` places with sign extension.

`fract32 shr_fr1x32(fract32 src, short shft)`
Arithmetically shifts the `src` variable right by `shft` places with sign extension. If `shft` is negative, the shift is to the left by `abs(shft)` places, and the empty bits are zero-filled.

`fract32 shr_fr1x32_clip(fract32 src, short shft)`
Arithmetically shifts the `src` variable right by `shft` (clipped to 6 bits) places with sign extension. If `shft` is negative, the shift is to the left by `abs(shft)` places, and the empty bits are zero-filled.

`fract16 sat_fr1x32(fract32 f1)`
If `f1>0x00007fff`, it returns `0x7fff`. If `f1<0xffff8000`, it returns `0x8000`. Otherwise, it returns the lower 16 bits of `f1`.

`fract16 round_fr1x32(fract32 f1)`
Rounds the 32-bit `fract` to a 16-bit `fract` using biased rounding.

`int norm_fr1x32(fract32)`
Returns the number of left shifts required to normalize the input variable so that it is either in the interval `0x40000000` to `0x7fffffff`, or in the interval `0x80000000` to `0xc0000000`. In other words,

```
fract32 x;
shl_fr1x32(x,norm_fr1x32(x));
```

returns a value in the range `0x40000000` to `0x7fffffff`, or in the range `0x80000000` to `0xc0000000`.

`fract16 trunc_fr1x32(fract32 f1)`
Returns the top 16 bits of `f1`—it truncates `f1` to 16 bits.

fract2x16 Built-in Functions

All built-in functions described here are saturating unless otherwise stated. These built-ins operate primarily on the `fract2x16` type, although there are composition and decomposition functions for the `fract2x16` type, multiplies that return `fract32` results, and operations on a single `fract2x16` pair that return `fract16` types.

The notation used to represent two `fract16` values packed into a `fract2x16` is `{a,b}`, where "`a`" is the `fract16` packed into the high half, and "`b`" is the `fract16` packed into the low half.

`fract2x16 compose_fr2x16(fract16 f1, fract16 f2)`

Takes two `fract16` values, and returns a `fract2x16` value.
Input: two `fract16` values
Returns: `{f1,f2}`

`fract16 high_of_fr2x16(fract2x16 f)`

Takes a `fract2x16` and returns the "high half" `fract16`.
Input: `f{a,b}`
Returns: `a`

`fract16 low_of_fr2x16(fract2x16 f)`

Takes a `fract2x16` and returns the "low half" `fract16`
Input: f{a,b}
Returns: b

```
fract2x16 add_fr2x16(fract2x16 f1,fract2x16 f2)
```

Adds two packed fracts.
Input: `f1{a,b} f2{c,d}`
Returns: `{a+c,b+d}`

```
fract2x16 sub_fr2x16(fract2x16 f1,fract2x16 f2)
```

Subtracts two packed fracts.
Input: `f1{a,b} f2{c,d}`
Returns: `{a-c,b-d}`

```
fract2x16 mult_fr2x16(fract2x16 f1,fract2x16 f2)
```

Multiplies two packed fracts. Truncates the results to 16 bits.
Input: `f1{a,b} f2{c,d}`
Returns: `{trunc16(a*c),trunc16(b*d)}`

```
fract2x16 multr_fr2x16(fract2x16 f1,fract2x16 f2)
```

Multiplies two packed fracts. Rounds the result to 16 bits. Whether the rounding is biased or unbiased depends on what the `RND_MOD` bit in the `ASTAT` register is set to.
Input: `f1{a,b} f2{c,d}`
Returns: `{round16{a*c},round16{b*d}}`

```
fract2x16 negate_fr2x16(fract2x16 f1)
```

Negates both 16-bit fracts in the packed `fract`. If one of the `fract16` values is `0x8000`, saturation occurs and `0x7fff` is the result of the negation.
Input: `f1{a,b}`
Returns: `{-a,-b}`

```
fract2x16 shl_fr2x16(fract2x16 f1,short shft)
```

Arithmetically shifts both `fract16s` in the `fract2x16` left by `shft` places, and returns the packed result. The empty bits are zero-filled. If `shft` is negative, the shift is to the right by `abs(shft)` places with sign extension.
Input: `f1{a,b} shft`
Returns: `{a<<shft,b<<shft}`

```
fract2x16 shl_fr2x16_clip(fract2x16 f1,short shft)
```

Arithmetically shifts both `fract16s` in the `fract2x16` left by `shft` (clipped to 5 bits) places, and returns the packed result. The empty bits are zero filled. If `shft` is negative, the shift is to the right by `abs(shft)` places with sign extension.

```
fract2x16 shr_fr2x16(fract2x16 f1,short shft)
```

Arithmetically shifts both `fract16s` in the `fract2x16` right by `shft` places with sign extension, and returns the packed result. If `shft` is negative, the shift is to the left by `abs(shft)` places and the empty bits are zero-filled.
Input: `f1{a,b} shft`
Returns: `{a>>shft,b>>shft}`

```
fract2x16 shr_fr2x16_clip(fract2x16 f1,short shft)
```

Arithmetically shifts both `fract16s` in the `fract2x16` right by `shft` (clipped to 5 bits) places with sign extension, and returns the packed result. If `shft` is negative, the shift is to the left by `abs(shft)` places and the empty bits are zero-filled.

```
fract2x16 shrl_fr2x16(fract2x16 f1,short shft)
```

Logically shifts both `fract16s` in the `fract2x16` right by `shft` places. There is no sign extension and no saturation – the empty bits are zero-filled.
Input: `f1{a,b} shft`
Returns: `{a>>shft,b>>shft}  //logical shift`

```
fract2x16 shrl_fr2x16_clip(fract2x16 f1,short shft)
```

Logically shifts both `fract16s` in the `fract2x16` right by `shft` places (clipped to 5 bits). There is no sign extension and no saturation – the empty bits are zero-filled.

```
fract2x16 abs_fr2x16(fract2x16 f1)
```

Returns the absolute value of both `fract16s` in the `fract2x16`.
Input: `f1{a,b}`
Returns: `{abs(a),abs(b)}`

```
fract2x16 min_fr2x16(fract2x16 f1,fract2x16 f2)
```

Returns the minimums of the two pairs of `fract16s` in the two input `fract2x16s`.
Input: `f1{a,b} f2{c,d}`
Returns: `{min(a,c),min(b,d)}`

```
fract2x16 max_fr2x16(fract2x16 f1,fract2x16 f2)
```

Returns the maximums of the two pairs of `fract16s` in the two input `fract2x16s`.
Input: `f1{a,b} f2{c,d}`
Returns: `{max(a,c),max(b,d)}`

```
fract16 sum_fr2x16(fract2x16 f1)
```

Performs a sideways addition of the two `fract16s` in `f1`.
Input: `f1{a,b}`
Returns: `(fract16) a+b`

```
fract2x16 add_as_fr2x16(fract2x16 f1,fract2x16 f2)
```

Performs a vector add/subtract on the two input `fract2x16s`.
Input: `f1{a,b} f2{c,d}`
Returns: `{a+c,b-d}`

```
fract2x16 add_sa_fr2x16(fract2x16 f1,fract2x16 f2)
```

Performs a vector subtract/add on the two input `fract2x16s`.
Input: `f1{a,b} f2{c,d}`
Returns: `{a-c,b+d}`

```
fract16 diff_hl_fr2x16(fract2x16 f1)
```

Takes the difference (high-low) of the two f`ract16s` in the `fract2x16`.
Input: `f1{a,b}`
Returns: `(fract16) a-b`

```
fract16 diff_lh_fr2x16(fract2x16 f1)
```

Takes the difference (low-high) of the two `fract16s` in the `fract2x16`.
Input: `f1{a,b}`
Returns: `(fract16) b-a`

```
fract32 mult_ll_fr2x16(fract2x16 f1, fract2x16 f2)
```

Cross-over multiplication. Multiplies the low half of `f1` with the low half of `f2`.
Input: `f1{a,b} f2{c,d}`
Returns: `(fract32) b*d`

```
fract32 mult_hl_fr2x16(fract2x16 f1, fract2x16 f2)
```

Cross-over multiplication. Multiplies the high half of `f1` with the low half of `f2`.
Input: `f1{a,b} f2{c,d}`
Returns: `(fract32) a*d`

```
fract32 mult_lh_fr2x16(fract2x16 f1, fract2x16 f2)
```

Cross-over multiplication. Multiplies the low half of `f1` with the high half of `f2`.
Input: `f1{a,b} f2{c,d}`
Returns: `(fract32) b*c`

```
fract32 mult_hh_fr2x16(fract2x16 f1, fract2x16 f2)
```

Cross-over multiplication. Multiplies the high half of `f1` with the high half of `f2`.
Input: `f1{a,b} f2{c,d}`
Returns: `(fract32) a*c`

ETSI Built-in Functions

If `fract.h` is included with `ETSI_SOURCE` defined, the macros listed below are also defined, mapping from the ETSI `fract` functions onto the compiler built-in functions. The mappings are done in `fract_math.h` (included by `fract.h`).

```
add()                    abs_s()
sub()                    saturate()
shl()                    extract_h()
shr()                    extract_l()
mult()                   L_deposit_l()
mult_r()                 div_s()
negate()                 norm_s()
round()                  norm_l()
L_add()                  L_Extract()
L_sub()                  L_Comp()
L_abs()                  Mpy_32()
L_negate()               Mpy_32_16()
L_shl()                  L_mult()
L_shr()                  L_mac()
L_msu()                  L_shr_r()
div_l()
```

Here is a description of the ETSI functions that do not map exactly to compiler built-in functions:

```
fract32 L_mac(fract32 acc,fract16 f1, fract16 f2)
```
Multiply accumulate. Returns `acc+=f1*f2`.

```
fract32 L_msu(fract32 acc,fract16 f1, fract16 f2)
```
Multiply subtract. Returns `acc-=f1*f2`.

`fract32 L_Comp(fract16 f1, fract16 f2)`
Returns `f1<<16 + f2<<1`.

`fract32 Mpy_32_16(short hi, short lo, fract16 n)`
Multiplies a `fract32` (decomposed to `hi` and `lo`) by a `fract16`, and returns the result as a `fract32`.

`void L_Extract(fract32 f1, fract16 *f2, fract16 *f3)`
Decomposes a 32-bit `fract` into two 16-bit `fracts`.

`fract32 Mpy_32(short hi1, short lo1, short hi2, short lo2)`
Multiplies two `fract32` numbers, and returns the result as a `fract32`. The input `fracts` have both been split up into two `shorts`.

`fract16 div_s(fract16 f1, fract16 f2)`
Produces a result which is the fractional division of `f1` by `f2`. Not a built-in function as written in C code.

By default, the ETSI functions

```
fract16 shl(fract16 _x, short _y);
fract16 shr(fract16 _x, short _y);
fract32 L_shl(fract32 _x, short _y);
fract32 L_shr(fract32 _x, short _y);
```

map to clipping versions of the built-in `fract` shifts. To map them to the faster, but non-clipping versions of the built-in fractional shifts, define the macro `_ADI_FAST_ETSI`, either in your source before you include `fract_math.h`, or on the compile command line.

ETSI Support

VisualDSP++ 4.0 for Blackfin processors provides European Telecommunications Standards Institute (ETSI) support routines in the `libetsi*.dlb` library. It contains routines for manipulation of the `fract16` and `fract32` data types as stipulated by ETSI. The routines provide bit-accurate calculations for common operations, and conversions between `fract16` and `fract32` data types.

To use the ETSI routines, the header file `libetsi.h` must be included, and all source code must be compiled with the `ETSI_SOURCE` macro defined.

These routines are:

- "32-Bit Fractional ETSI Routines" on page 1-119
- "16-Bit Fractional ETSI Routines" on page 1-123

Several of the ETSI routines are provided with carry and overflow checking. Where overflow or carry occurs, the global variables `Carry` and `Overflow` are set. It is your responsibility to reset these variables in between operations.

The `Carry` and `Overflow` variables are represented by integers and are prototyped in the `libetsi.h` system header file.

There are two types of `libetsi` libraries provided with VisualDSP++ 4.0. Those with a name of the form `libetsi*co.dlb` have been compiled with checking and setting of `Overflow` and `Carry`. Those with a name of the form `libetsi*.dlb` (with no "`co`") have the checking and setting of `Overflow` and `Carry` disabled for optimal performance. To use the `Carry` and `Overflow` checking versions of the library, use the compiler flag "`-l etsi*co`". When rebuilding `libetsi`, `Carry` and `Overflow` checking is enabled with the C and assembler macro definition `__SET_ETSI_FLAGS=1`

The carry/overflow setting functions libraries (`libetsi532co.dlb` and `libetsi535co.dlb`) are not built by default by the supplied makefiles. To rebuild the carry and overflow setting versions of the libraries, define compiler macro `__SET_ETSI_FLAGS =1` during compilation.

The carry/overflow setting versions of the following functions will not set the `Carry` and/or `Overflow` variables correctly on the BF535 processor, due to differences in the way the hardware flags are set on the BF535 processor.

```
shl       shr     shr_r
L_msuNs   L_shl   L_shr   L_shr_r
```

Many of the routines in the library are also represented by built-in functions. Where built-in functions exist, the compiler replaces the functional code with an optimal inline assembler representation. To disable the use of the ETSI built-in functions and use the library versions, compile with the macro `NO_ETSI_BUILTINS` defined. However, use of the built-in functions results in better performance since there is an overhead in making the function call to the library.

The built-in versions of the functions do not set the `Carry` and `Overflow` flags.

The built-in versions of some ETSI functions are affected by the `RND_MOD` flag in the `ASTAT` register. For bit-exact results, the `RND_MOD` flag should be set to provide biased rounding.

If the macro `RENAME_ETSI_NEGATE` is defined, the ETSI function "`negate`" will be renamed to "`etsi_negate`". This is useful because the C++ standard declares a template function called `negate` (found in the C++ include "`functional`").

The following routines are available in the ETSI library. These routines are commonly classified into two groups: those that return or primarily operate on 32-bit fractional data, and those that return or primarily operate on 16-bit fractional data.

32-Bit Fractional ETSI Routines

The following functions return or primarily operate on 32-bit fractional data.

`fract32 L_add_c(fract32, fract32)`

Performs 32-bit addition of the two input parameters. When using a version of the library compiled with `__SET_ETSI_FLAGS`, the `Carry` and `Overflow` flags are set when carry and overflow/underflow occur during addition.

`fract32 L_abs(fract32)`

Returns the 32-bit absolute value of the input parameter. In cases where the input is equal to `0x80000000`, saturation occurs and `0x7fffffff` is returned. A built-in version of this function is also provided.

`fract32 L_add(fract32, fract32)`

Returns the 32-bit saturated result of the addition of the two input parameters. If the library is compiled with `__SET_ETSI_FLAGS`, the `Overflow` flag are set when overflow occurs. A built-in version of this function is also provided.

`fract32 L_deposit_h(fract16)`

Deposits the 16-bit parameter into the 16 most significant bits of the 32-bit result. The least 16 bits are zeroed. A built-in version of this function is also provided.

`fract32 L_deposit_l(fract16)`

Deposits the 16-bit parameter into the 16 least significant bits of the 32-bit result. The most significant bits are set to sign extension for the input. A built-in version of this function is also provided.

`fract32 L_mac(fract32 acc, fract16 f1, fract16 f2)`

Performs a fractional multiplication of the two 16-bit parameters and returns the saturated sum of the multiplication result with the 32-bit parameter. A built-in version of this function is also provided.

`fract32 L_macNs(fract32,fract16, fract16)`

Performs a non-saturating version of the `L_mac` operation. If the library is compiled with `__SET_ETSI_FLAGS`, the `Overflow` and `Carry` flags are set when carry or overflow/underflow occurs.

`fract32 L_mls (fract32, fract16)`

Multiplies both the most significant bits and the least significant bits of a `long`, by the same `short`.

`fract32 L_msu(fract32, fract16, fract16)`

Performs a fractional multiplication of the two 16-bit parameters and returns the saturated subtraction of the multiplication result with the 32-bit parameter. A built-in version of this function is also provided.

`fract32 L_msuNs(fract32, fract16, fract16)`

Performs a non-saturating version of the `L_msu` operation. If the library is compiled with `__SET_ETSI_FLAGS`, the `Overflow` and `Carry` flags are set when carry or overflow/underflow occurs.

`fract32 L_mult(fract16, fract16)`

Returns the 32-bit saturated result of the fractional multiplication of the two 16-bit parameters. A built-in version of this function is also provided.

`fract32 L_negate(fract32)`

Returns the 32-bit result of the negation of the parameter. Where the input parameter is `0x80000000` saturation occurs and `0x7fffffff` is returned. A built-in version of this function is also provided.

`fract32 L_sat(fract32)`

The resultant variable is set to `0x80000000` if `Carry` and `Overflow` flags are set (underflow condition); else, if `Overflow` is set, the resultant is set to `0x7fffffff`. The default revision of the library simply returns as no checking or setting of the `Overflow` and `Carry` flags is performed.

`fract32 L_shl(fract32 src, fract16 shft)`

Arithmetically shifts the 32-bit first parameter to the left by the value given in the 16-bit second parameter. The empty bits of the 32-bit value are zero-filled. If the shifting value, `shft`, is negative, the source is shifted to the right by `-shft`, sign-extended. The result is saturated in cases of overflow and underflow.

If the library is compiled with `__SET_ETSI_FLAGS`, the `Overflow` flag is set when overflow occurs. A built-in version of this function is also provided.

`fract32 L_shr(fract32, fract16)`

Arithmetically shifts the 32-bit first parameter to the right by the value given in the 16-bit second parameter with sign extension. If the shifting value is negative, the source is shifted to the left. The result is saturated in cases of overflow and underflow.

If the library is compiled with `__SET_ETSI_FLAGS`, the `Overflow` flag is set when overflow occurs. A built-in version of this function is also provided.

`fract32 L_shr_r(fract32, fract16)`

Performs the shift-right operation as per `L_shr` but with rounding. If the library is compiled with `__SET_ETSI_FLAGS`, the `Overflow` and `Carry` flags are set when carry or overflow/underflow occurs.

`fract32 L_sub(fract32, fract32)`

Returns the 32-bit saturated result of the subtraction of two 32-bit parameters (first-second). A built-in version of this function is also provided.

`fract32 L_sub_c(fract32 f1, fract32 f2)`

Performs 32-bit subtraction of two fractional values (`f1 - f2`). When using a version of the library compiled with `__SET_ETSI_FLAGS`, the `Carry` and `Overflow` flags are set when carry and overflow/underflow occur during subtraction.

`fract32 L_Comp(fract16 hi,fract16 lo)`

Composes a `fract32` type value from the given `fract16` high and low components. The sign is provided with the low half, and the result is calculated as:

```
hi<<16 + lo<<1;
```

A built-in version of this function is also provided.

`fract32 Mpy_32(fract16 hi1,fract16 lo1,fract16 hi2,fract16 lo2)`

Performs the multiplication of two `fract32` type variables, provided as high and low half parameters. The result returned is calculated as:

```
Res = L_mult(hi1,hi2);
Res = L_mac(Res, mult(hi1,lo2),1);
Res = L_mac(Res, mult(lo1,hi2),1);
```

A built-in version of this function is also provided.

```
fract32 Mpy_32_16(fract16 hi,fract16 lo,fract16 v)
```

Multiplies a `fract16` type value by a `fract32` type value provided as high and low halves, and return the result as a `fract32`. A built-in version of this function is also provided.

```
void L_Extract(fract32 src,fract16 *hi,fract16 *lo)
```

Extracts low and high halves of `fract32` type value into `fract16` variables pointed to by the high and low parameters. The values calculated are:

```
Hi = bit16 to bit31 of src
Lo = (src - hi<<16)>>1
```

A built-in version of this function is also provided.

```
fract32 Div_32(fract32 L_num,fract16 denom_hi,fract16 denom_lo)
```

Performs 32-bit fractional divide operation.
The result returned is the `fract32` representation of `L_num` divided by `L_denom` (represented by `demon_hi` and `denom_lo`). Both `L_num` and `L_denom` must be positive fractional values and `L_num` must be less than `L_denom` to ensure that the result falls within the fractional range.

16-Bit Fractional ETSI Routines

The following functions return or primarily operate on 16-bit fractional data.

```
fract16 abs_s(fract16)
```

Returns the 16-bit value that is the absolute value of the input parameter. Where the input is `0x8000`, saturation occurs and `0x7fff` is returned. A built-in version of this function is also provided.

`fract16 add(fract16, fract16)`

Returns the 16-bit result of adding the two `fract16` input parameters.

Saturation occurs with the result being set to `0x7fff` for overflow and `0x8000` for underflow. If the library is compiled with `__SET_ETSI_FLAGS`, the `Overflow` and `Carry` flags are set when carry or overflow/underflow occurs. A built-in version of this function is also provided.

`fract16 div_l (fract32, fract16)`

This function produces a result which is the fractional integer division of the first parameter by the second. Both inputs must be positive and the least significant word of the second parameter must be greater or equal to the first; the result is positive (leading bit equal to 0) and truncated to 16 bits. The function calls `abort()` on division error conditions.

`fract16 div_s(fract16 f1, fract16 f2)`

Returns the 16-bit result of the fractional integer division of `f1` by `f2`. Both `f1` and `f2` must be positive fractional values with `f2` greater than `f1`. A built-in version of this function is also provided.

`fract16 extract_l(fract32)`

Returns the 16 least significant bits of the 32-bit fract parameter provided. A built-in version of this function is also available.

`fract16 extract_h(fract32)`

Returns the 16 most significant bits of the 32-bit fract parameter provided. A built-in version of this function is also available.

`fract16 mac_r(fract32 acc, fract16 f1, fract16 f2)`

Performs an `L_mac` operation using the three parameters provided. The result is the rounded 16 most significant bits of the 32-bit results from the `L_mac` operation.

`fract16 msu_r(fract32, fract16, fract16)`

Performs an `L_msu` operation using the three parameters provided. The result is the rounded 16 most significant bits of the 32-bit result from the `L_msu` operation.

`fract16 mult(fract16, fract16)`

Returns the 16-bit result of the fractional multiplication of the input parameters. The result is saturated. A built-in version of this function is also provided.

`fract16 mult_r(fract16, fract16)`

Performs a 16-bit multiply with rounding of the result of the fractional multiplication of the two input parameters. A built-in version of this function is also provided.

The inline version of the `mult_r()` function is implemented using the `multr_fr1x16()` compiler intrinsic, which in turn does a normal 16-bit fractional multiply:

```
Rx.L = Ry.L * Rz.L;
```

This instruction's result is affected by the `RND_MOD` bit in the ASTAT register, which means that the results are not always ETSI-compliant. To avoid this issue, set `RND_MOD` before using the inline version or use the `libetsi` library-defined version of the function (which sets the bit).

`fract16 negate(fract16)`

Returns the 16-bit result of the negation of the input parameter. If the input is `0x8000`, saturation occurs and `0x7fff` is returned. A built-in version of this function is also provided.

`int norm_l(fract32)`

Returns the number of left shifts required to normalize the input variable so that it is either in the interval `0x40000000` to `0x7fffffff`, or in the interval `0x80000000` to `0xc0000000`. In other words,

```
fract32 x;
shl_fr1x32(x,norm_fr1x32(x));
```

returns a value in the range `0x40000000` to `0x7fffffff`, or in the range `0x80000000` to `0xc0000000`.

`int norm_s(fract16)`

Returns the number of left shifts required to normalize the input variable so that it is either in the interval `0x4000` to `0x7fff`, or in the interval `0x8000` to `0xc000`. In other words,

```
fract16 x;
shl_fr1x16(x,norm_fr1x16(x));
```

returns a value in the range `0x4000` to `0x7fff`, or in the range `0x8000` to `0xc000`.

`fract16 round(fract32)`

Rounds the lower 16 bits of the 32-bit input parameter into the most significant 16 bits with saturation. The resulting bits are shifted right by 16. A built-in version of this function is also provided.

`fract16 saturate(fract32)`

Return the 16 most significant bits of the input parameter. If the input parameter is greater than `0x7fff`, `0x7fff` is returned. If the input parameter is less than `0x8000`, `0x8000` is returned. A built-in version of this function is also available.

`fract16 shl(fract16 src, fract16 shft)`

Arithmetically shifts the `src` variable left by `shft` places. The empty bits are zero-filled. If `shft` is negative, the shift is to the right by `shft` places.

If the library is compiled with `__SET_ETSI_FLAGS`, the `Overflow` and `Carry` flags are set when carry or overflow/underflow occurs. A built-in version of this function is also provided.

`fract16 shr(fract16, fract16)`

Arithmetically shifts the `src` variable right by `shft` places with sign extension. If `shft` is negative, the shift is to the left by `shft` places.

If the library is compiled with `__SET_ETSI_FLAGS`, the `Overflow` and `Carry` flags are set when carry or overflow/underflow occurs. A built-in version of this function is also provided.

`fract16 shr_r(fract16, fract16)`

Performs a shift to the right as per the `shr()` operation with additional rounding and saturation of the result.

`fract16 sub(fract16 f1, fract16 f2)`

Returns the 16-bit result of the subtraction of the two parameters (`f1 - f2`). Saturation occurs with the result being set to `0x7fff` for overflow and `0x8000` for underflow.

If the library is compiled with `__SET_ETSI_FLAGS`, the `Overflow` and `Carry` flags are set when carry or overflow/underflow occurs. A built-in version of this function is also provided.

Fractional Value Built-In Functions in C++

The compiler provides support for two C++ fractional classes. The `fract` class uses a `fract32` C type for storage of the fractional value, whereas `shortfract` uses a `fract16` C type for storage of the fractional value.

Instances of the `shortfract` and `fract` class can be initialized using values with the "`r`" suffix, provided they are within the range `[-1,1)`. The `fract` class is represented by the compiler as representing the internal type `fract`. For example,

```
#include <fract>
int main ()
{
   fract X = 0.5r;
}
```

Instances of the `shortfract` class can be initialized using "`r`" values in the same way, but are not represented as an internal type by the compiler. Instead, the compiler produces a temporary `fract`, which is initialized using the "`r`" value. The value of the `fract` class is then copied to the `shortfract` class using an implicit copy and the `fract` is destroyed.

The `fract` and `shortfract` classes contain routines that allow basic arithmetic operations and movement of data to and from other data types. The example below shows the use of the `shortfract` class with * and + operators.

The mathematical routines for addition, subtraction, division and multiplication for both `fract` and `shortfract` classes are performed using the ETSI-defined routines for the C fractional types (`fract16` and `fract32`). Inclusion of the `fract` and `shortfract` header files implicitly defines the macro `ETSI_SOURCE` to be 1. This is required for use of the ETSI routines, which are defined in `libetsi.h` and located in the `libetsi53*.dlb` libraries.

```
#include <shortfract>
#include <stdio.h>
```

```
#define N 20

shortfract x[N] = {
   .5r,.5r,.5r,.5r,.5r,
   .5r,.5r,.5r,.5r,.5r,
   .5r,.5r,.5r,.5r,.5r,
.5r,.5r,.5r,.5r,.5r};

shortfract y[N] = {
   0,.1r,.2r,.3r,.4r,
   .5r,.6r,.7r,.8r,.9r,
   .10r,.1r,.2r,.3r,.4r,
   .5r,.6r,.7r,.8r,.9r};

shortfract fdot(int n, shortfract *x, shortfract *y)
{
   int j;
   shortfract s;
   s = 0;
   for (j=0; j<n; j++) {
         s += x[j] * y[j];
   }
   return s;
}

int main(void)
{
   fdot(N,x,y);
}
```

Fractional Literal Values in C

The "r" suffix is not available when compiling in C mode, since "r" literals are instances of the `fract` class. However, if a C program is compiled in C++ mode, `fract16` and `fract32` variables can be initialized using "r" literal values; the compiler automatically converts from the class values to the C types. When adopting this approach, be aware of any semantic differences between the C and C++ languages that might affect your program.

Complex Fractional Built-In Functions in C

The `complex_fract16` type is used to hold complex fractional numbers. It contains real and imaginary values, both as 16-bit fractional numbers.

```
typedef struct {
   fract16 re, im;
} complex_fract16;
```

The `complex_fract32` type is used to hold complex fractional numbers. It contains real and imaginary values, both as 32-bit fractional numbers.

```
typedef struct {
   fract32 re, im;
} complex_fract32;
```

The `complex_fract16` and `complex_fract32` types are defined by the `complex.h` header file. Additionally, there are numerous library functions for manipulating complex fracts. These functions are documented in "DSP Run-Time Library Reference" on page 4-43.

The compiler also supports the following built-in operations for complex `fracts`.

- Complex fractional multiply and accumulate and multiply and subtract

  ```
  cmac_fr16();
  cmsu_fr16();
  ```

- Complex fractional square

  ```
  csqu_fr16();
  csqu_fr32();
  ```

- Complex fractional distance

  ```
  cdst_fr16();
  cdst_fr32();
  ```

Complex Operations in C++

The C++ complex class is defined in the `<complex>` header file, and defines a template class for manipulating complex numbers. The standard arithmetic operators are overloaded, and there are `real()` and `imag()` methods for obtaining the relevant part of the complex number.

For example, the determinate and inverse of a `2x2` matrix of complex doubles may be computed using the following C++ function:

```
#include <complex>

complex<double> inverse2d(const complex<double> mx[4],
complex<double> mxinv[4])
{
   complex<double> det = mx[0] * mx[3] - mx[2] * mx[1];

   if( (det.real() != 0.0) || (det.imag() != 0.0) ) {
      complex<double> invdet = complex<double>(1.0,0.0) / det;

      mxinv[0] =   invdet * mx[3];
      mxinv[1] = -(invdet * mx[1]);
      mxinv[2] = -(invdet * mx[2]);
      mxinv[3] =   invdet * mx[0];
   }
   return det;
}
```

As a comparison, the equivalent function in C is:

```
#include <complex.h>

complex_double inverse2d(const complex_double mx[4],
complex_double mxinv[4])
{
   complex_double det;
   complex_double invdet;
   complex_double tmp;

   det = cmlt(mx[0],mx[3]);
   tmp = cmlt(mx[2],mx[1]);
```

```
    det = csub(det,tmp);

    if( (det.re != 0.0) || (det.im != 0.0) ) {
       invdet = cdiv((complex_double){1.0,0.0},det);

       mxinv[0] = cmlt(invdet,mx[3]);
       mxinv[1] = cmlt(invdet,mx[1]);
       mxinv[1].re = -mxinv[1].re;
       mxinv[1].im = -mxinv[1].im;
       mxinv[2] = cmlt(invdet,mx[2]);
       mxinv[2].re = -mxinv[2].re;
       mxinv[2].im = -mxinv[2].im;
       mxinv[3] = cmlt(invdet,mx[0]);
    }
    return det;
}
```

Viterbi History and Decoding Functions

There are four built-in functions available which provide the selection function of a Viterbi decoder. Specifically, these four functions provide the maximum value selection and history update parts. The functions all make use of the `A0` accumulator to maintain the history value. (The accumulator register maintains the history values by shifting the previous value along one place and setting a bit to indicate the result of the current iteration's selection.)

You must include the `ccblkfn.h` file before Viterbi functions are used. Failure to do so leads to unresolved symbols at link time.

The four Viterbi functions allow for left or right shifting (setting the least or most significant bit, accordingly) and for `1x16` or `2x16` operands.

The four Viterbi functions are multi-valued; they update some of their parameters in place, since Viterbi operations return both the selection result and the updated history. The first two Viterbi functions provide left- and right-shifting operations for single 16-bit input operands.

The functions are:

```
void lvitmax1x16(int value, int oldhist,
     short selected, int newhist)
void rvitmax1x16(int value, int oldhist,
     short selected, int newhist)
```

The two functions, `lvitmax1x16()` and `rvitmax1x16()`, perform "selection-and-update" operations for two 16-bit operands, which are in the high and low halves of "`value`", respectively. The "`oldhist`" operand contains the history value from the preceding iteration. The "`selected`" and "`newhist`" operands are not inputs at all; instead, their expressions must be `lvalues` (valid on the left-hand side of an assignment), whose values are updated by the operation.

The "`selected`" operand is set to contain the largest half of value. The "`newhist`" operand is set to contain the `oldhist` value, shifted one place (left for `lvitmax`, right for `rvitmax`), and with one bit (LSB for `lvitmax`, MSB for `rvitmax`) set to 1 if the high half was selected; 0 otherwise.

The second two Viterbi functions provide left- and right-shifting operations for pairs of 16-bit input operands. The functions are:

```
lvitmax2x16(int value_x, int value_y, int oldhist,
      short selected, int newhist)
void rvitmax2x16(int value_x, int value_y, int oldhist,
      short selected, int newhist)
```

The two functions, `lvitmax2x16()` and `rvitmax2x16()`, perform two selection-and-update operations. Each of the `value_x` and `value_y` input expressions contains two 16-bit operands. A selection operation is performed on the two 16-bit operands in `value_x`, and another selection operation is performed on the two 16-bit operands in `value_y`. The `oldhist` value is shifted and updated into `newhist`, as described above.

However, in this example, `oldhist` is shifted two places, and two bits are set. The history value is shifted one place, and a bit is set to indicate the result of the `value_x` selection operation. Then, the history value is shifted a second place, and another bit is set to indicate the result for the `value_y` selection operation.

The selected value from `value_x` is stored in the low half of `selected`. The selected value from `value_y` is stored in the high half of `selected`.

Circular Buffer Built-In Functions

The C/C++ compiler provides the built-in functions that use the Blackfin processor's circular buffer mechanisms. These functions provide automatic circular buffer generation, circular indexing, and circular pointer references.

Automatic Circular Buffer Generation

If optimization is enabled, the compiler automatically attempts to use circular buffer mechanisms where appropriate. For example,

```
void func(int *array,int n,int incr)
{
int i;
for (i = 0;i < n;i++)
    array [ i % 10 ] += incr;
}
```

The compiler recognizes that the "`[i % 10 ]`" expression is a circular reference, and uses a circular buffer if possible. There are cases where the compiler is unable to verify that the memory access is always within the bounds of the buffer. The compiler is conservative in such cases, and does not generate circular buffer accesses.

The compiler can be instructed to still generate circular buffer accesses even in such cases, by specifying the `-force-circbuf` switch (on page 1-32).

Explicit Circular Buffer Generation

The compiler also provides built-in functions that can explicitly generate circular buffer accesses, subject to available hardware resources. The built-in functions provide circular indexing and circular pointer references. Both built-in functions are defined in the `ccblkfn.h` header file.

Circular Buffer Increment of an Index

The following operation performs a circular buffer increment of an index.

```
long circindex(long index, long incr, unsigned long nitems);
```

The operation is equivalent to:

```
index += incr;
if (index < 0)
   index += nitems;
else if (index >= nitems)
   index -= nitems;
```

An example of this built-in function is:

```
#include <ccblkfn.h>
void func(int *array, int n, int incr, int len)
{
      int i, idx = 0;

      for (i = 0; i < n; i++) {
         array[idx] += incr;
         idx = circindex(idx, incr, len);
      }
}
```

Circular Buffer Increment of a Pointer

The following operation performs a circular buffer increment of a pointer.

```
void *circptr(void *ptr, long incr , void * base,  unsigned long buflen);
```

Both *incr* and *buflen* are specified in bytes, since the operation deals in void pointers.

The operation is equivalent to:

```
ptr += incr;
if (ptr < base)
    ptr += buflen;
else if (ptr >= (base+buflen))
    ptr -= buflen;
```

An example of this built-in function is:

```
#include <ccblkfn.h>
void func(int *array, int n, int incr, int len)
{
      int i, idx = 0;
      int *ptr = array;

      // scale increment and length by size
      // of item pointed to.
      incr *= sizeof(*ptr);
      len *= sizeof(*ptr);

      for (i = 0; i < n; i++) {
           *ptr += incr;
           ptr = circptr(ptr, incr, array, len);
      }
}
```

Endian-Swapping Intrinsics

The following two intrinsics are available for changing data from big-endian to little-endian, or vice versa.

```
#include <ccblkfn.h>
int byteswap4(int);
short byteswap2(short);
```

For example, `byteswap2(0x1234)` returns `0x3412`.

Since Blackfin processors use a little-endian architecture, these intrinsics are useful when communicating with big-endian devices, or when using a protocol that requires big-endian format. For example,

```
struct bige_buffer {
        int len;
        char data[MAXLEN];
} buf;

int i, len;
buf = get_next_buffer();
len = byteswap4(buf.len);
for (i = 0; i < len; i++)
        process_byte(buf.data[i]);
```

System Built-In Functions

The following built-in functions allow access to system facilities on Blackfin processors. The functions are defined in the `ccblkfn.h` header file. Include the `ccblkfn.h` file before using these functions. Failure to do so leads to unresolved symbols at link time.

Stack Space Allocation

```
void *alloca(unsigned)
```

This function allocates the requested number of bytes on the local stack, and returns a pointer to the start of the buffer. The space is freed when the current function exits. The compiler supports this function via `__builtin_alloca()`.

System Register Values

```
unsigned int sysreg_read(int reg)
void sysreg_write(int reg, unsigned int val)
unsigned long long sysreg_read64(int reg)
void sysreg_write64(int reg,unsigned long long val)
```

These functions get (read) or set (write) the value of a system register. In all cases, `reg` is a constant from the file `<sysreg.h>`.

IMASK Values

```
unsigned cli(void)
void sti(unsigned mask)
```

The `cli()` function retrieves the old value of `IMASK`, and disables interrupts by setting `IMASK` to all zeros. The `sti()` function installs a new value into `IMASK`, enabling the interrupt system according to the new mask stored.

Interrupts and Exceptions

```
void raise_intr(int)
void excpt(int)
```

These two functions raise interrupts and exceptions, respectively. In both cases, the parameter supplied must be an integer literal value.

Idle Mode

```
void idle(void)
```

places the processor in idle mode.

Synchronization

```
void csync(void)
void ssync(void)
```

These two functions provide synchronization. The `csync()` function is a core-only synchronization—it flushes the pipeline and store buffers. The `ssync()` function is a system synchronization, and also waits for an `ACK` instruction from the system bus.

Compiler Performance Built-in Functions

These functions provide the compiler with information about the expected behavior of the program. You can use these built-in functions to tell the compiler which parts of the program are most likely to be executed; the compiler can then arrange for the most common cases to be those that execute most efficiently.

```
#include <ccblkfn.h>
int expected_true(int cond);
int expected_false(int cond);
```

For example, consider the code

```
extern int func(int);
int example(int call_the_function, int value)
{
   int r = 0;
   if (call_the_function)
      r = func(value);
   return r;
}
```

If you expect that parameter `call_the_function` to be true in the majority of cases, you can write the function in the following manner:

```
extern int func(int);
int example(int call_the_function, int value)
{
   int r = 0;
   if (expected_true(call_the_function))
                    // indicate most likely true
      r = func(value);
   return r;
}
```

This indicates to the compiler that you expect `call_the_function` to be true in most cases, so the compiler arranges for the default case to be to call function `func()`. If, on the other hand, you were to write the function as:

```
extern int func(int);
int example(int call_the_function, int value)
{
   int r = 0;
   if (expected_false(call_the_function))
                     // indicate most likely false
      r = func(value);
   return r;
}
```

then the compiler arranges for the generated code to default to the opposite case, of not calling function `func()`.

These built-in functions do not change the operation of the generated code, which will still evaluate the boolean expression as normal. Instead, they indicate to the compiler which flow of control is most likely, helping the compiler to ensure that the most commonly-executed path is the one that uses the most efficient instruction sequence.

The `expected_true` and `expected_false` built-in functions only take effect when optimization is enabled in the compiler. They are only supported in conditional expressions.

Video Operation Built-In Functions

The C/C++ compiler provides built-in functions for using the Blackfin processor's video pixel operations. Include the `video.h` header file before using these functions.

Some video operation built-in functions take an 8-byte sequence of data, and select from it a sequence of four bytes to use as input. The operation selects the four bytes at an offset of 0, 1, 2, or 3 bytes from lowest byte of the 8-byte sequence, depending on the value of a pointer parameter.

Where reverse variants of the operations exist (the operation name is suffixed by "r"), the two 4-byte halves of the 8-byte sequence are accessed in reverse order.

Where a video operation generates more than one result, the operation may be implemented by more than one built-in function. In these cases, macros are provided to generate the appropriate built-in calls.

For further information regarding the underlying Blackfin processor instructions that implement the video operations, refer to Chapter 13 in the *Instruction Set Reference* for the appropriate target processor.

Function Prototypes

Align Operations

```
int align8(int src1, int src2);    /* 1 byte offset */
int align16(int src1, int src2);   /* 2 byte offset */
int align24(int src1, int src2);   /* 3 byte offset */
```

These three operations treat their two inputs as a single 8-byte sequence, and extract a specific 4-byte sequence from it, starting at offset 1, 2 or 3 bytes, as shown.

Packing Operations

```
int bytepack(int src1, int src2);
```

This operation treats its two inputs as four 16-bit values, and packs each 16-bit value into an 8-bit value in the result. Effectively, it converts an array of four `shorts` to an array of four `chars`.

```
long long compose_i64(int low, int high);
```

This operation produces a 64-bit value from the two 32-bit values provided as input and can be used to efficiently generate a `long long` type that is needed for many of the following operations.

Disaligned Loads

```
int loadbytes(int *ptr);
```

This operation is used to load a 4-byte sequence from memory using `ptr` as the address, where `ptr` may be misaligned. The actual data retrieved is aligned, by masking off the bottom two bits of `ptr`, where `ptr` is intended to select bytes from input operands in subsequent operations. Misaligned read exceptions are prevented from occurring.

Unpacking

```
byteunpack(long long src, char *ptr, int dst1, int dst2)

byteunpackr(long long src, char *ptr, int dst1, int dst2)
```

These macros provide the unpacking operations, where `PTR` selects four bytes from the eight-byte sequence in `SRC`. Each of the four bytes is expanded to a 16-bit value. The first two 16-bit values are returned in `DST1`, and the second two are returned in `DST2`.

Quad 8-Bit Add Subtract

```
add_i4x8(long long src1, char *ptr1, long long src2,
     char *ptr2, int dst1, int dst2);

add_i4x8r(long long src1, char *ptr1, long long src2,
     char *ptr2, int dst1, int dst2);

sub_i4x8(long long src1, char *ptr1, long long src2,
     char *ptr2, int dst1, int dst2);

sub_i4x8r(long long src1, char *ptr1, long long src2,
     char *ptr2, int dst1, int dst2);
```

These macros provide the operations to select two four-byte sequences from the two eight-byte operands provided, either add or subtract the corresponding bytes, and generate four 16-bit results. The first two results are stored in `DST1`, and the second two are stored in `DST2`. `PTR1` selects the bytes from `SRC1`, and `PTR2` selects the bytes from `SRC2`. The `add_i4x8r()`

and sub_i4x8r() variants produce the same instructions as add_i4x8() and sub_i4x8(), but with the "reverse" option enabled; this swaps the order of the two 32-bit elements in the SRC parameters.

Dual 16-Bit Add/Clip

```
int addclip_lo(long long src1, char *ptr1, long long src2,
    char *ptr2);

int addclip_hi(long long src1, char *ptr1, long long src2,
    char *ptr2);

int addclip_lor(long long src1, char *ptr1, long long src2,
    char *ptr2);

int addclip_hir(long long src1, char *ptr1, long long src2,
    char *ptr2);
```

These operations select two 16-bit values from src1 using ptr1, and two 8-bit values from src2 using ptr2. The pairs are added and then clipped to the range 0 to 255, producing two 8-bit results. The _lo versions select bytes 3 and 1 from src2, while the _hi versions select bytes 2 and 0. The _lor and _hir versions reverse the order of the 32-bit elements in src1 and src2.

Quad 8-Bit Average

```
int avg_i4x8(long long src1, char *ptr1, long long src2,
    char *ptr2);

int avg_i4x8_t(long long src1, char *ptr1, long long src2,
    char *ptr2);

int avg_i4x8_r(long long src1, char *ptr1, long long src2,
    char *ptr2);

int avg_i4x8_tr(long long src1, char *ptr1, long long src2,
    char *ptr2);
```

These operations select two 4-byte sequences from `src1` and `src2`, using `ptr1` and `ptr2`. They add the corresponding bytes from each sequence, and then shift each result right once to produce four byte-size averages. There are four variants of the operation to select the reverse and truncate options for the operation.

```
int avg_i2x8_lo(long long src1, char *ptr1, long long src2);
int avg_i2x8_lot(long long src1, char *ptr1, long long src2);
int avg_i2x8_lor(long long src1, char *ptr1, long long src2);
int avg_i2x8_lotr(long long src1, char *ptr1, long long src2);
int avg_i2x8_hi(long long src1, char *ptr1, long long src2);
int avg_i2x8_hit(long long src1, char *ptr1, long long src2);
int avg_i2x8_hir(long long src1, char *ptr1, long long src2);
int avg_i2x8_hitr(long long src1, char *ptr1, long long src2);
```

These operations all produce two 8-bit average values. Each selects two four-byte sequences from `src1` and `src2`, using `ptr` and then produces averages of the 4-byte sequences as two `2x2` byte clusters. The two results are byte-sized, and are stored in two bytes of the output result; the other two bytes are set to zero. The variants allow for the generation of different options: truncate or round, reverse input pairs, or store results in the low or high bytes of each 16-bit half of the result register.

Accumulator Extract With Addition

```
extract_and_add(long long src1, long long src2, int dst1,
      int dst2);
```

This macro provides the operation to add the high and low halves of `SRC1` with the high and low halves of `SRC2` to produce two 32-bit results.

Subtract Absolute Accumulate

```
saa(long long src1, char *ptr1, long long src2, char *ptr2,
    int sum1, int sum2, int dst1, int dst2);

saar(long long src1, char *ptr1, long long src2, char *ptr2,
    int sum1, int sum2, int dst1, int dst2);
```

These macros provide the operations to select two 4-byte sequences from `SRC1` and `SRC2`, using `PTR1` and `PTR2` to select. The bytes from `SRC2` are subtracted from their corresponding bytes in `SRC1`, and then the absolute value of each subtraction is computed. These four results are then added to the four 16-bit values in `SUM1` and `SUM2`, and the results are stored in `DST1` and `DST2`, as four 16-bit values.

Example of Use: Sum of Absolute Difference

As an example use of the video operation built-in functions, a block-based video motion estimation algorithm might use Sum of Absolute Difference (SAD) calculations to measure distortion. A reference SAD function may be implemented as:

```
int ref_SAD16x16(unsigned char *image, unsigned char *block,
   int imgWidth)
{
   int dist = 0;
   int x, y;
   for (y = 0; y < 16; y++) {
      for (x = 0; x < 16; x++)
         dist += abs(image[x] - block[x]);
      image += 16;
      block += 16;
   }
   return dist;
}
```

Using video operation built-in functions, the code could be written as:

```
int vid_SAD16x16(char *image, char *block,
   int imgWidth)
{
   int dist = 0;
   int x, y;
   long long srcI, srcB;
   char *ptrI, *ptrB;
   int bytesI1, bytesI2, bytesB1, bytesB2;
   int sum1, sum2, res1, res2;
   sum1 = sum2 = 0;
   bytesI2 = bytesB2 = 0;
   for (y = 0; y < 16; y++) {
        ptrI = image;
        ptrB = block;
        for (x = 0; x < 16; x += 8) {
            bytesI1 = loadbytes((int *)ptrI); ptrI += 4;
            bytesB1 = loadbytes((int *)ptrB); ptrB += 4;
            srcI = compose_i64(bytesI1, bytesI2);
            srcB = compose_i64(bytesB1, bytesB2);
            saa(srcI, ptrI, srcB, ptrB, sum1, sum2, sum1, sum2);

            bytesI2 = loadbytes((int *)ptrI); ptrI += 4;
            bytesB2 = loadbytes((int *)ptrB); ptrB += 4;
            srcI = compose_i64(bytesI1, bytesI2);
            srcB = compose_i64(bytesB1, bytesB2);
            saar(srcI, ptrI, srcB, ptrB, sum1, sum2, sum1, sum2);
        }
   }
   extract_and_add(sum1, sum2, res1, res2);
   dist = res1 + res2;
   return dist;
}
```

Misaligned Data Built-In Functions

The following intrinsic functions exist to allow the user to explicitly perform loads from misaligned memory locations and stores to misaligned memory locations. These functions always generate expanded code to read and write from such memory locations, regardless of whether the access is aligned or not.

```
#include <ccblkfn.h>

short misaligned_load16(void *);
short misaligned_load16_vol(volatile void *);
void misaligned_store16(void *, short);
void misaligned_store16_vol(volatile void *, short);

int misaligned_load32(void *);
int misaligned_load32_vol(volatile void *);
void misaligned_store32(void *, int);
void misaligned_store32_vol(volatile void *, int);

long long misaligned_load64(void *);
long long misaligned_load64_vol(volatile void *);
void misaligned_store64(void *, long long);
void misaligned_store64_vol(volatile void *, long long);
```

Note that there are also volatile variants of these functions. Because of the operations required to read from and write to such misaligned memory locations, no assumptions should be made regarding the atomicity of these operations. Refer to "#pragma pack (alignopt)" on page 1-153 for more information.

Pragmas

The Blackfin C/C++ compiler supports a number of pragmas. Pragmas are implementation-specific directives that modify the compiler's behavior. There are two types of pragma usage: pragma directives and pragma operators.

Pragma directives have the following syntax:

```
#pragma pragma-directive pragma-directive-operands new-line
```

Pragma operators have the following syntax:

```
_Pragma ( string-literal )
```

When processing a pragma operator, the compiler effectively turns it into a pragma directive using a non-string version of *string-literal*. This means that the following pragma directive

```
#pragma linkage_name mylinkname
```

can also be equivalently expressed using the following pragma operator,

```
_Pragma ("linkage_name mylinkname")
```

The examples in this manual use the directive form.

The C compiler supports pragmas for:

- Arranging alignment of data
- Defining functions that can act as interrupt handlers
- Changing the optimization level, midway through a module
- Changing how an externally visible function is linked
- Providing header file configurations and properties
- Giving additional information about loop usage to improve optimizations

The following sections describe the supported pragmas.

- "Data Alignment Pragmas" on page 1-149
- "Interrupt Handler Pragmas" on page 1-156
- "Loop Optimization Pragmas" on page 1-156
- "General Optimization Pragmas" on page 1-159
- "Linking Control Pragmas" on page 1-160
- "Function Side-Effect Pragmas" on page 1-167
- "Template Instantiation Pragmas" on page 1-173
- "Header File Control Pragmas" on page 1-175
- "Diagnostic Control Pragmas" on page 1-178
- "Memory Bank Pragmas" on page 1-181

The compiler issues a warning when it encounters an unrecognized pragma directive or pragma operator. The compiler does not expand any preprocessor macros used within any pragma directive or pragma operator.

Data Alignment Pragmas

The data alignment pragmas are used to modify how the compiler arranges data within the processor's memory. Since the Blackfin processor architecture requires memory accesses to be naturally aligned, each data item is normally aligned at least as strongly as itself—two-byte `shorts` have an alignment of 2, and four-byte `longs` have an alignment of 4. The 8-byte `long longs` also have an alignment of 4.

When `structs` are defined, the struct's overall alignment is the same as the field which has the largest alignment. The struct's size may need padding to ensure all fields are properly aligned and that the struct's overall size is a multiple of its alignment.

Sometimes, it is useful to change these alignments. A `struct` may have its alignment increased to improve the compiler's opportunities in vectorizing access to the data. A `struct` may have its alignment reduced so that a large array occupies less space.

If a data item's alignment is reduced, the compiler cannot safely access the data item without the risk of causing misaligned memory access exceptions. Programs that use reduced-alignment data must ensure that accesses to the data are made using data types that match the reduced alignment, rather than the default one. For example, if an `int` has its alignment reduced from the default (4) to 2, it must be accessed as two `shorts` or four bytes, rather than as a single `int`.

The data alignment pragmas include `align`, `pack` and `pad` pragmas. Alignments specified using these pragmas must be a power of two. The compiler rejects uses of those pragmas that specify alignments that are not powers of two.

#pragma align *num*

This pragma may be used before variable declarations and field declarations. It applies to the variable or field declaration that immediately follows the pragma.

The pragma's effect is that the next variable or field declaration should be forced to be aligned on a boundary specified by *num*.

- If *num* is greater than the alignment normally required by the following variable or field declaration, then the variable or field declaration's alignment is changed to *num*.
- If *num* is less than the alignment normally required, then the variable or field declaration's alignment is changed to *num*, and a warning is given that the alignment has been reduced.

The pragma also allows the following keywords as allowable alignment specifications:

- `_WORD` – Specifies a 32-bit alignment
- `_LONG` – Specifies a 64-bit alignment
- `_QUAD` – Specifies a 128-bit alignment

If the `pack` or `pad` pragmas (see in this section) are currently active, then `align` overrides the immediately following field declaration. The following are the examples of how to use `#pragma align`.

```
struct s{
#pragma align 8    /* field a aligned on 8-byte boundary    */
   int a;
   int bar;

#pragma align 16   /* field b aligned on 16-byte boundary   */
   int b;
} t[2];

#pragma align 256
int arr[128];      /* declares an int array with 256 alignment */
```

The following example shows a use that is valid, but causes a compiler warning.

```
#pragma align 1
int warns;         /* declares an int with byte alignment,   */
                   /* causes a compiler warning              */
```

The following is an example of an invalid use of `#pragma align`; because the alignment is not a power of two, the compiler rejects it and issues an error.

```
#pragma align 3
int errs;      /* INVALID: declares an int with non-power of */
               /* two alignment, causes a compiler error     */
```

The `align` pragma only applies to the immediately-following definition, even if that definition is part of a list. For example,

```
#pragma align 8
int i1, i2, i3;       // pragma only applies to i1
```

#pragma alignment_region (*alignopt*)

Sometimes it is desirable to specify an alignment for a number of consecutive data items rather than individually. This can be done using the `alignment_region` and `alignment_region_end` pragmas:

- `#pragma alignment_region` sets the alignment for all following data symbols up to the corresponding `alignment_region_end` pragma
- `#pragma alignment_region_end` removes the effect of the active alignment region and restores the default alignment rules for data symbols.

The rules concerning the argument are the same as for `#pragma align`. The compiler faults an invalid alignment (such as an alignment that is not a power of two). The compiler warns if the alignment of a data symbol within the control of an `alignment_region` is reduced below its natural alignment (as for `#pragma align`).

Use of the `align` pragma overrides the region alignment specified by the currently active `alignment_region` pragma (if there is one). The currently active `alignment_region` does not affect the alignment of fields.

Example:

```
#pragma align 16

int aa;          /* alignment 16 */
int bb;          /* alignment 4  */

#pragma alignment_region (8)
```

```
int cc;          /* alignment 8 */
int dd;          /* alignment 8 */
int ee;          /* alignment 8 */

#pragma align 16

int ff;          /* alignment 16 */
int gg;          /* alignment 8  */
int hh;          /* alignment 8  */

#pragma alignment_region_end

int ii;          /* alignment 4 */

#pragma alignment_region (2)

long double jj;   /* alignment 2, but the compiler warns
                     about the reduction */

#pragma alignment_region_end

#pragma alignment_region (5)
long double kk;   /* the compiler faults this, alignment is not
                     a power of two */

#pragma alignment_region_end
```

#pragma pack (*alignopt*)

This pragma may be applied to `struct` definitions. It applies to all `struct` definitions that follow, until the default alignment is restored by omitting *`alignopt`*; for example, by `#pragma pack()` with empty parentheses.

The pragma is used to reduce the default alignment of the `struct` to be aligned. If there are fields within the `struct` that have a default alignment greater than `align`, their alignment is reduced to `alignopt`. If there are fields within the struct that have alignment less than `align`, their alignment is unchanged.

If `alignopt` is specified, it is illegal to invoke `#pragma pad` until the default alignment is restored. The compiler generates an error message if the `pad` and `pack` pragmas are used in a manner that conflicts.

The following shows how to use `#pragma pack`:

```
#pragma pack(1)
/* struct minimum alignment now 1 byte, uses of
   "#pragma pad" would cause a compilation error now */

struct is_packed {
   char a;
   /* normally the compiler would add three padding bytes here,
   but not now because of prior pragma pack use */
   int b;
} t[2];            /* t definition requires 10 packed bytes */

#pragma pack()
/* struct minimum alignment now not one byte,
   "#pragma pad"can now be used legally */

struct is_packed u[2];   /* u definition requires 10 packed
                            bytes */
/* struct not_packed is a new type, and will not be packed.*/

struct not_packed {
   char a;
   /* compiler will insert three padding bytes here */
   int b;
} w[2];                  /* w definition required 16 bytes */
```

On architectures that do not support misaligned memory accesses at the hardware level, the compiler generates additional code to correctly handle reads from (and writes to) misaligned structure members. The code generated will not necessarily be as efficient as reading from (or writing to) an aligned structure member, but that is the trade-off that must be accepted in return for getting packed structures.

Only direct reads from (and writes to) misaligned structure members are automatically handled by the compiler. As a result, taking the address of a misaligned field and assigning it to a pointer causes the compiler to generate a warning. The reason for the warning is that the compiler does not detect a misaligned memory access if the address of a misaligned field is taken and stored in a pointer of a different type to that of the structure.

#pragma pad (*alignopt*)

This pragma may be applied to `struct` definitions. It applies to `struct` definitions that follow until the default alignment is restored by omitting *`alignopt`*; for example, by `#pragma pad()` with empty parentheses.

This pragma is effectively shorthand for placing `#pragma align` before every field within the `struct` definition. Like pragma `pack`, it reduces the alignment of fields which default to an alignment greater than *`alignopt`*.

However, unlike pragma `pack`, it also increases the alignment of fields that default to an alignment less than *`alignopt`*.

While pragma `pack` *`alignopt`* generates a warning if a field alignment is reduced, pragma `pad` *`alignopt`* does not.

If *`alignopt`* is specified, it is illegal to invoke `#pragma pack` until the default alignment is restored.

The following example shows how to use `#pragma pad()`.

```
#pragma pad(4)
struct {
   int i;
   int j;
} s = {1,2};
#pragma pad()
```

Interrupt Handler Pragmas

The `interrupt`, `nmi`, and `exception` pragmas all declare that the following function declaration or definition is to be used as an entry in the Event Vector Table (EVT). The compiler arranges for the function to save its context. This is more than the usual called-preserved set of registers. The function returns using an instruction appropriate to the type of event specified by the pragma.

These pragmas are not normally used directly; there are macros provided by the `<sys/exception.h>` file. See "Interrupt Handler Support" on page 1-201 for more information.

The pragmas may be specified on either the function's declaration or its definition. Only one of the three pragmas listed above may be specified for a particular function.

The `interrupt_reentrant` pragma is used with the `interrupt` pragma to specify that the function's context-saving prologue should also arrange for interrupts to be re-enabled for the duration of the function's execution.

Loop Optimization Pragmas

Loop optimization pragmas give the compiler additional information about usage within a particular loop, which allows the compiler to perform more aggressive optimization. The pragmas are placed before the loop statement, and apply to the statement that immediately follows, which must be a `for`, `while` or `do` statement to have effect. In general, it is most effective to apply loop pragmas to inner-most loops, since the compiler can achieve the most savings there.

The optimizer always attempts to vectorize loops when it is safe to do so. The optimizer exploits the information generated by the interprocedural analysis (see "Interprocedural Analysis" on page 1-74) to increase the cases where it knows it is safe to do so.

Consider the code:

```
void copy(short *a, short *b) {
     int i;
     for (i=0; i<100; i++)
         a[i] = b[i];
}
```

If you call `copy` with two calls, such as `copy(x,y)` and later `copy(y,z)`, the interprocedural analysis is unable to tell that "`a`" never aliases "`b`". Therefore, the optimizer cannot be sure that one iteration of the loop is not dependent on the data calculated by the previous iteration of the loop. If it is known that each iteration of the loop is not dependent on the previous iteration, then the `vector_for` pragma can be used to explicitly notify the compiler that this is the case.

#pragma all_aligned

This pragma applies to the subsequent loop. This pragma asserts that all pointers are initially aligned on the most desirable boundary.

#pragma different_banks

This pragma allows the compiler to assume that groups of memory accesses based on different pointers within a loop reside in different memory banks. By scheduling them together, memory access is much improved.

#pragma loop_count(*min, max, modulo*)

This pragma appears just before the loop it describes. It asserts that the loop iterates at least `min` times, no more than `max` times, and a multiple of `modulo` times. This information enables the optimizer to omit loop guards and to decide whether the loop is worth completely unrolling and whether code needs to be generated for odd iterations. The last two arguments can be omitted if they are unknown.

For example,

```
int i;
#pragma loop_count(24, 48, 8)
for (i=0; i < n; i++)
```

#pragma no_alias

Use this pragma to tell the compiler the following loop has no loads or stores that conflict. When the compiler finds memory accesses that potentially refer to the same location through different pointers (known as "aliases"), the compiler is restricted in how it may reorder or vectorize the loop, because all the accesses from earlier iterations must be complete before the compiler can arrange for the next iteration to start.

For example,

```
void vadd(int *a, int *b, int *out, int n) {
   int i;
#pragma no_alias
   for (i=0; i < n; i++)
       out[i] = a[i] + b[i];
}
```

The `no_alias` pragma appears just before the loop it describes. This pragma asserts that in the next loop, no `load` or `store` operation conflict with each other. In other words, no `load` or `store` in any iteration of the loop has the same address as any other `load` or `store` in the current or in any other iteration of the loop. In the example above, if the pointers `a` and `b` point to two memory areas that do not overlap, then no load from `b` is using the same address as any store to `a`. Therefore, `a` is never an alias for `b`.

Using the `no_alias` pragma can lead to better code because it allows any number of iterations to be performed concurrently (rather than just two at a time), thus providing better software pipelining by the optimizer.

#pragma no_vectorization

This pragma turns off all vectorization for the loop on which it is specified.

#pragma vector_for

This pragma notifies the optimizer that it is safe to execute two iterations of the loop in parallel. The `vector_for` pragma does not force the compiler to vectorize the loop. The optimizer checks various properties of the loop and does not vectorize it if it believes to be unsafe or if it cannot deduce that the various properties necessary for the vectorization transformation are valid.

Strictly speaking, the pragma simply disables checking for loop-carried dependencies.

```
void copy(short *a, short *b) {
int i;
#pragma vector_for
   for (i=0; i<100; i++)
a[i] = b[i];
}
```

In cases where vectorization is impossible (for example, if array `a` were aligned on a word boundary but array `b` was not), the information given in the assertion made by `vector_for` may still be put to good use in aiding other optimizations.

General Optimization Pragmas

The compiler supports several pragmas which can change the optimization level while a given module is being compiled. These pragmas must be used globally, immediately prior to a function definition. The pragmas do not just apply to the immediately following function; they remain in effect until the end of the compilation, or until superseded by one of the following `optimize_` pragmas.

- **#pragma optimize_off**
  This pragma turns off the optimizer, if it was enabled. This pragma has no effect if Interprocedural Optimization Analysis is enabled.

- **#pragma optimize_for_space**
  This pragma turns the optimizer back on, if it was disabled, or sets the focus to give reduced code size a higher priority than high performance, where these conflict.

- **#pragma optimize_for_speed**
  This pragma turns the optimizer back on, if it was disabled, or sets the focus to give high performance a higher priority than reduced code size, where these conflict.

- **#pragma optimize_as_cmd_line**
  This pragma resets the optimization settings to be those specified on the `ccblkfn` command line when the compiler was invoked.

These are code examples for the `optimize_` pragmas.

```
#pragma optimize_off
void non_op() { /* non-optimized code */ }

#pragma optimize_for_space
void op_for_si() { /* code optimized for size */ }

#pragma optimize_for_speed
void op_for_sp() { /* code optimized for speed */ }
/* subsequent functions declarations optimized for speed */
```

Linking Control Pragmas

Linking pragmas (`linkage_name`, `retain_name`, `section` and `weak_entry`) change how a given global function or variable is viewed during the linking stage.

#pragma linkage_name *identifier*

This pragma associates the *identifier* with the next external function declaration. It ensures that the *identifier* is used as the external reference, instead of following the compiler's usual conventions. If the *identifier* is not a valid function name, as could be used in normal function definitions, the compiler generates an error. See also the `asm` keyword (described on page 1-194).

The following shows an example use of this pragma.

```
#pragma linkage_name realfuncname
void funcname ();
void func() {
   funcname();   /* compiler will generate a call to realfuncname */
}
```

#pragma retain_name

This pragma indicates that the function or variable declaration that follows the pragma is not removed even though IPA sees that it is not used. Use this pragma for C functions that are called only from an assembler routine, such as the startup code sequence invoked before `main()`.

The following example shows how to use this pragma.

```
int delete_me(int x) {
   return x-2;
}

#pragma retain_name
int keep_me(int y) {
   return y+2;
}

int main(void) {
   return 0;
}
```

Since the program has no uses for either `delete_me()` or `keep_me()`, the compiler removes `delete_me()`, but keeps `keep_me()` because of the pragma. You do not need to specify `retain_name` for `main()`.

For more information on IPA, see "Interprocedural Analysis" on page 1-74.

#pragma section/#pragma default_section

The section pragmas provide greater control over the sections in which the compiler places symbols.

The `section(`*`SECTSTRING`* `[, `*`QUALIFIER`*`, ...])` pragma is used to override the target section for any global or static symbol immediately following it. The pragma allows greater control over section qualifiers compared to the `section` keyword.

The `default_section(`*`SECTKIND`* `[, `*`SECTSTRING`* `[, `*`QUALIFIER`*`, ...]])` pragma is used to override the default sections in which the compiler is placing its symbols. The default sections fall into five different categories (listed under *`SECTKIND`*), and this pragma remains in force for a section category until its next use with that particular category. The omission of a section name results in the default section being reset to be the section that was in use at the start of processing.

`SECTKIND` can be one of the following keywords:

`CODE`	Section is used to contain procedures and functions
`ALLDATA`	Section is used to contain any data (normal, read-only and uninitialized)
`DATA`	Section is used to contain "normal data"
`CONSTDATA`	Section is used to contain read-only data
`BSZ`	Section is used to contain uninitialized data

`SECTSTRING` is the double-quoted string containing the section name, exactly as it will appear in the assembler file.

`QUALIFIER` can be one of the following keywords:

ZERO_INIT	Section will be zero-initialized at program startup
NO_INIT	Section will not be initialized at program startup
RUNTIME_INIT	Section will be user-initialized at program startup
DOUBLE32	Section may contain 32-bit but not 64-bit doubles
DOUBLE64	Section may contain 64-bit but not 32-bit doubles
DOUBLEANY	Section may contain either 32-bit or 64-bit doubles

There may be any number of comma-separated section qualifiers within such pragmas, but they must not conflict with one another. Qualifiers must also be consistent across pragmas for identical section names, and omission of qualifiers is not allowed even if at least one such qualifier has appeared in a previous pragma for the same section. If any qualifiers have not been specified for a particular section by the end of the translation unit, the compiler uses default qualifiers appropriate for the target processor. The compiler always tries to honor the `section` pragma as its highest priority, and the `default_section` pragma is always the lowest priority.

For example, the following code results in function `f` being placed in the section `foo`:

```
#pragma default_section(CODE, "bar")
#pragma section("foo")
void f() {}
```

The following code results in `x` being placed in section `zeromem`:

```
#pragma default_section(BSZ, "zeromem")
int x;
```

#pragma symbolic_ref

This pragma may be used before a public global variable, to indicate to the compiler that references to that variable should only be through the variable's symbolic name. Loading the address of a variable into a pointer register can be an expensive operation, and the compiler usually avoids this when possible. Consider the case where

```
int x;
int y;
int z;
void foo(void) { x = y + z; }
```

Given that the three variables are in the same data section, the compiler can generate the following code:

```
_foo:
            P0.L = .epcbss;
            P0.H = .epcbss;
            R0 = [P0+ 4];
            R1 = [P0+ 8];
            R0 = R1 + R0;
            [P0+ 0] = R0;
            RTS;

      .section/ZERO_INIT bsz;

      .align 4;
.epcbss:
      .type .epcbss,STT_OBJECT;
     .byte _x[4];
     .global _x;
     .type _x,STT_OBJECT;
     .byte _y[4];
     .global _y;
     .type _y,STT_OBJECT;
     .byte _z[4];
     .global _z;
     .type _z,STT_OBJECT;
.epcbss.end:
```

Having loaded a pointer to "x" (which shares the address of the start of the .epcbss section), the compiler can use offsets from this pointer to access "y" and "z", avoiding the expense of loading addresses for those variables. However, this forces the linker to ensure that the relative offsets between x, y, z and .epcbss do not change during the linking process.

There are cases when you might wish the compiler to reference a variable only through its symbolic name, such as when you are using RESOLVE() in the .LDF file to explicitly map the variable to a particular address. The compiler automatically uses symbolic references for:

- volatile variables
- variables with #pragma weak_entry specified
- variables greater or equal to 16 bytes in size

If other cases arise, you can use #pragma symbolic_ref to explicitly request this behavior. For example,

```
int x;
#pragma symbolic_ref
int y;
int z;
void foo(void) { x = y + z; }
```

produces

```
_foo:
            P0.L = .epcbss;
            I0.L = _y;
            P0.H = .epcbss;
            I0.H = _y;
            MNOP || R0 = [P0+ 4] || R1 = [I0];
            R0 = R0 + R1;
            [P0+ 0] = R0;
            RTS;

     .section/ZERO_INIT bsz;
```

```
        .align 4;
    .epcbss:
        .type .epcbss,STT_OBJECT;
        .byte _x[4];
        .global _x;
        .type _x,STT_OBJECT;
        .byte _z[4];
        .global _z;
        .type _z,STT_OBJECT;
    .epcbss.end:
        .align 4;
        .global _y;
        .type _y,STT_OBJECT;
        .byte _y[4];
    ._y.end:
```

Note that variable `y` is referenced explicitly by name, rather than using the common pointer to `.epcbss`, and it is declared outside the bounds of the (`.epcbss, .epcbss.end`) pair. The (`_y, ._y.end`) form a separate pair that can be moved by the linker if necessary without affecting the functionality of the generated code.

The `symbolic_ref` pragma can only be used immediately before declarations of global variables, and only applies to the immediately-following declaration.

#pragma weak_entry

This pragma may be used before a static variable or function declaration or definition. It applies to the function/variable declaration or definition that immediately follows the pragma. Use of this pragma causes the compiler to generate the function or variable definition with weak linkage.

The following are example uses of the `pragma weak_entry` directive.

```
#pragma weak_entry
int w_var = 0;

#pragma weak_entry
void w_func(){}
```

Function Side-Effect Pragmas

The function side-effect pragmas (`alloc`, `pure`, `const`, `regs_clobbered`, and `result_alignment`) are used before a function declaration to give the compiler additional information about the function in order to improve the code surrounding the function call. These pragmas should be placed before a function declaration and should apply to that function. For example,

```
#pragma pure
long dot(short*, short*, int);
```

#pragma alloc

This pragma tells the compiler that the function behaves like the library function "`malloc`", returning a pointer to a newly allocated object. An important property of these functions is that the pointer returned by the function does not point at any other object in the context of the call. In the example,

```
#define N 100

#pragma alloc
int *new_buf(void);
int *vmul(int *a, int *b) {
   int *out = new_buf();
   for (i = 0; i < N; ++i)
      out[i] = a[i] * b[i];
   return out;
}
```

the compiler can reorder the iterations of the loop because the `#pragma alloc` tells it that `a` and `b` cannot overlap `out`.

The GNU attribute `malloc` is also supported with the same meaning.

#pragma const

This pragma is a more restrictive form of the `pure` pragma. It tells the compiler that the function does not read from global variables as well as not write to them or read or write volatile variables. The result is therefore a function of its parameters. If any of the parameters are pointers, the function may not read the data they point at.

#pragma noreturn

This pragma can be placed before a function prototype or definition. Its use tells the compiler that the function to which it applies will never return to its caller. For example, a function such as the standard C function "exit" never returns.

The use of this pragma allows the compiler to treat all code following a call to a function declared with the pragma as unreachable and hence removable.

```
#pragma noreturn
void func() {
   while(1);
}

main() {
   func();
   /* any code here will be removed */
}
```

#pragma pure

This pragma tells the compiler that the function does not write to any global variables, and does not read or write any volatile variables. Its result, therefore, is a function of its parameters or of global variables. If any of the parameters are pointers, the function may read the data they point at but it may not write to the data.

Since the function call has the same effect every time it is called (between assignments to global variables), the compiler need not generate the code for every call. Therefore, in this example,

```
#pragma pure
long sdot(short *, short *, int);

long tendots(short *a, short *b, int n) {
   int i;
   long s = 0;
for (i = 1; i < 10; ++i)
      s += sdot(a, b, n);  // call can get hoisted out of loop
   return s;}
```

the compiler can replace the ten calls to `sdot` with a single call made before the loop.

#pragma regs_clobbered *string*

This pragma may be used with a function declaration or definition to specify which registers are modified (or clobbered) by that function. The *`string`* contains a list of registers and is case-insensitive.

When used with an external function declaration, this pragma acts as an assertion telling the compiler something it would not be able to discover for itself. In the example,

```
#pragma regs_clobbered "r5 p5 i3"
void f(void);
```

the compiler knows that only registers `r5`, `p5` and `i3` may be modified by the call to `f`, so it may keep local variables in other registers across that call.

The `regs_clobbered` pragma may also be used with a function definition, or a declaration preceding a definition (when it acts as a command to the compiler to generate register saves, and restores on entry and exit from the function) to ensure it only modifies the registers in string.

For example,

```
#pragma regs_clobbered "r3 m4 p5"
int g(int a) {
   return a+3;
}
```

The `regs_clobbered` pragma may not be used in conjunction with `#pragma interrupt`. If both are specified, a warning is issued and the `regs_clobbered` pragma is ignored.

To obtain optimum results with the pragma, it is best to restrict the clobbered set to be a subset of the default scratch registers. When considering when to apply the `regs_clobbered` pragma, it may be useful to look at the output of the compiler to see how many scratch registers were used. Restricting the volatile set to these registers will produce no impact on the code produced for the function but may free up registers for the caller to allocate across the call site.

The `regs_clobbered` pragma cannot be used in any way with pointers to functions. A function pointer cannot be declared to have a customized clobber set, and it cannot take the address of a function which has a customized clobber set. The compiler raises an error if either of these actions are attempted.

String Syntax

A `regs_clobbered` string consists of a list of registers, register ranges, or register sets that are clobbered (Table 1-12). The list is separated by spaces, commas, or semicolons.

A *`register`* is a single register name—the same name may be used in an assembly file.

A *`register range`* consists of `start` and `end` registers which both reside in the same register class, separated by a hyphen. All registers between the two (inclusive) are clobbered.

A *register set* is a name for a specific set of commonly-clobbered registers that is predefined by the compiler. Table 1-12 shows defined register sets,

Table 1-12. Clobbered Register Sets

Set	Registers
CCset	ASTAT, condition codes
DAGscratch	Members of I, M, B and L registers that are scratch by default
Dscratch	Members of D registers that are scratch by default, ASTAT
Pscratch	Members of P registers that are scratch by default
DPscratch	Dscratch union Pscratch
ALLscratch	Entire default volatile set

When the compiler detects an illegal string, a warning is issued and the default volatile set (see in "Scratch Registers" on page 1-252) is used instead.

Unclobberable and Must Clobber Registers
There are certain caveats as to what registers may or must be placed in the clobbered set (Table 1-12).

On Blackfin processors, the registers SP and FP may not be specified in the clobbered set, as the correct operation of the function call requires their values to be preserved. If the user specifies them in the clobbered set, a warning is issued and they are removed from the specified clobbered set.

Registers from these classes,

```
I, P, D, M, ASTAT, A0, A1, LC, LT, LB
```

may be specified in the clobbered set and code is generated to save them as necessary.

The L registers are required to be zero on entry and exit from a function. A user may specify that a function clobbers the L registers. If it is a compiler-generated function, then it leaves the L registers zero at the end of the function. If it is an assembly function, then it may clobber the L registers. In that case, the L registers are re-zeroed after any call to that function.

The `SEQSTAT`, `RETI`, `RETX`, `RETN`, `SYSCFG`, `CYCLES` and `CYCLES2` registers are never used by the compiler and are never preserved.

Register `P1` is used by the linker to expand `CALL` instructions, so it may be modified at the call site regardless of whether the `regs_clobbered` pragma says it is clobbered or not. Therefore, the compiler never keeps `P1` live across a call. However, the compiler accepts the pragma when compiling a function in case the user wants to keep `P1` live across a call that is not expanded by the linker. It is your responsibility to make sure such calls are not expanded by the linker.

User Reserved Registers
User reserved registers, which are indicated via the `-reserve` switch (on page 1-53), are never be preserved in the function wrappers whether in the clobbered set or not.

Function Parameters
Function calling conventions are visible to the caller and do not affect the clobbered set that may be used on a function. For example,

```
#pragma regs_clobbered ""   // clobbers nothing
void f(int a, int b);
void g() {
   f(2,3);
}
```

The parameters `a` and `b` are passed in registers `R0` and `R1`, respectively. No matter what happens in function `f`, after the call returns, the values of `R0` and `R1` would still be 2 and 3, respectively.

Function Results
The registers in which a function returns its result must always be clobbered by the callee and retain their new value in the caller. They may appear in the clobbered set of the callee but it does not make any difference to the generated code—the return register are not saved and restored. Only the return register used by the particular function return type is special. Return registers used by different return types are treated in the clobbered list in the convention way.

For example,

```
typedef struct { int x, int y } Point;
typedef struct { int x[10] } Big;
int f();   // Result in R0. R1,P0 may be preserved across call.
Point g(); // Result in R0 and R1. P0 may be preserved across call.
Big f();   // Result pointer in P0. R0, R1 may be preserved
              accross call.
```

#pragma result_alignment (*n*)

This pragma asserts that the pointer or integer returned by the function has a value that is a multiple of *n*.The pragma is often used in conjunction with the `#pragma alloc` of custom-allocation functions that return pointers more strictly aligned than could be deduced from their type.

Template Instantiation Pragmas

The template instantiation pragmas (`instantiate`, `do_not_instantiate` and `can_instantiate`) give fine grain control over where (that is, in which object file) the individual instances of template functions, member functions, and static members of template classes are created. The creation of these instances from a template is known in "C++ speak" as instantiation. As templates are a feature of C++, these pragmas are allowed only in `-c++` mode.

Refer to "Compiler C++ Template Support" on page 1-275 for more information on how the compiler handles templates.

The instantiation pragmas take the name of an instance as a parameter, as shown in Table 1-13.

Table 1-13. Instance Names

Name	Parameter
a template class name	`A<int>`
a template class declaration	`class A<int>`
a member function name	`A<int>::f`
a static data member name	`A<int>::I`
a static data declaration	`int A<int>::I`
a member function declaration	`void A<int>::f(int, char)`
a template function declaration	`char* f(int, float)`

If the instantiation pragmas are not used, the compiler selects object files where all required instances automatically instantiate during the prelinking process.

#pragma instantiate *instance*

This pragma requests the compiler to instantiate *`instance`* in the current compilation. For example,

```
#pragma instantiate class Stack<int>
```

causes all static members and member functions for the `int` instance of a template class `Stack` to be instantiated, whether they are required in this compilation or not. The example,

```
#pragma instantiate void Stack<int>::push(int)
```

causes only the individual member function `Stack<int>::push(int)` to be instantiated.

#pragma do_not_instantiate *instance*

This pragma directs the compiler not to instantiate *instance* in the current compilation. For example,

```
#pragma do_not_instantiate int Stack<float>::use_count
```

prevents the compiler from instantiating the static data member `Stack<float>::use_count` in the current compilation.

#pragma can_instantiate *instance*

This pragma tells the compiler that if instance is required anywhere in the program, it should be instantiated in this compilation.

Currently, this pragma forces the instantiation even if it is not required anywhere in the program. Therefore, it has the same effect as `#pragma instantiate`.

Header File Control Pragmas

The header file control pragmas (`hdrstop`, `no_implicit_inclusion`, `no_pch`, `once`, and `system_header`) help the compiler to handle header files.

#pragma hdrstop

This pragma is used with the `-pch` (precompiled header) switch (on page 1-49). The switch instructs the compiler to look for a precompiled header (`.pch` file), and, if it cannot find one, to generate a file for use on a later compilation. The `.pch` file contains a snapshot of all the code preceding the header stop point.

By default, the header stop point is the first non-preprocessing token in the primary source file. The `#pragma hdrstop` can be used to set the point earlier in the source file.

In the example,

```
#include "standard_defs.h"
#include "common_data.h"
#include "frequently_changing_data.h"

int i;
```

the default header stop point is the start of the declaration `i`. This might not be a good choice, as in this example, "`frequently_changing_data.h`" might change frequently, causing the `.pch` file to be regenerated often, and, therefore, losing the benefit of precompiled headers. The `hdrstop` pragma can be used to move the header stop to a more appropriate place.

For the following example,

```
#include "standard_defs.h"
#include "common_data.h"
#pragma hdrstop
#include "frequently_changing_data.h"

int i;
```

the precompiled header file would not include the contents of `frequently_changing_data.h`, as it is included after the `hdrstop` pragma, and so the precompiled header file would not need to be regenerated each time `frequently_changing_data.h` was modified.

#pragma no_implicit_inclusion

With the `-c++` switch (on page 1-24), for each included `.h` file, the compiler attempts to include the corresponding `.c` or `.cpp` file. This is called implicit inclusion.

If `#pragma no_implicit_inclusion` is placed in an `.h` file, the compiler does not implicitly include the corresponding `.c` or `.cpp` file with the `-c++` switch. This behavior only affects the `.h` file with `#pragma no_implicit_inclusion` within it and the corresponding `.c` or `.cpp` files.

For example, if there are the following files,

`t.c` containing

```
#include “m.h”
```

and `m.h` and `m.c` are both empty, then

```
ccblkfn -c++ t.c -M
```

shows the following dependencies for `t.c`:

```
t.doj: t.c
t.doj: m.h
t.doj: m.C
```

If the following line is added to `m.h`,

```
#pragma no_implicit_inclusion
```

running the compiler as before would not show `m.c` in the dependencies list, such as:

```
t.doj: t.c
t.doj: m.h
```

#pragma no_pch

This pragma overrides the `-pch` (precompiled headers) switch (on page 1-49) for a particular source file. It directs the compiler not to look for a `.pch` file and not to generate one for the specified source file.

#pragma once

This pragma, which should appear at the beginning of a header file, tells the compiler that the header is written in such a way that including it several times has the same effect as including it once. For example,

```
#pragma once
#ifndef FILE_H
#define FILE_H
```

```
... contents of header file ...
#endif
```

In this example, the `#pragma once` is actually optional because the compiler recognizes the `#ifndef/#define/#endif` idiom and does not reopen a header that uses it.

#pragma system_header

This pragma identifies an `include` file as the file supplied with VisualDSP++.

The pragma tells the compiler that every function and variable declared in the file (but not in files included in the file) is the variable or function with that name from the VisualDSP++ library. The compiler takes advantage of any special knowledge it has of the behavior of the library.

Diagnostic Control Pragmas

The compiler supports `#pragma diag(`*`action: diag`* `[, `*`diag`* `...])` which allows selective modification of the severity of compiler diagnostic messages.

The directive has three forms:

- modify the severity of specific diagnostics
- modify the behavior of an entire class of diagnostics
- save or restore the current behavior of all diagnostics

Modifying the Severity of Specific Diagnostics

This form of the directive has the following syntax:

```
#pragma diag(ACTION: DIAG [, DIAG ...])
```

The *`action:`* qualifier can be one of the following keywords:

Keyword	Action
`suppress`	Suppresses all instances of the diagnostic
`remark`	Changes the severity of the diagnostic to a remark.
`warning`	Changes the severity of the diagnostic to a warning.
`error`	Changes the severity of the diagnostic to an error.
`restore`	Restores the severity of the diagnostic to what it was originally at the start of compilation after all command-line options were processed.

The *`diag`* qualifier can be one or more comma-separated compiler diagnostic numbers without the preceding "`cc`" (but can include leading zeros). The choice of error numbers is limited to those that may have their severity overridden (such as those that are displayed with a "`{D}`" in the error message). In addition, those diagnostics that are emitted by the compiler back-end (for example, after lexical analysis and parsing) cannot have their severity overridden either. Any attempt to override diagnostics that may not have their severity changed is silently ignored.

Modifying the Behavior of an Entire Class of Diagnostics

This form of the directive has the following syntax:

```
#pragma diag(ACTION)
```

The effects are as follows:

- **#pragma diag*(errors)***

  This pragma can be used to inhibit all subsequent warnings and remarks (equivalent to the `-w` switch option).

- **#pragma diag*(remarks)***

  This pragma can be used to enable all subsequent remarks and warnings (equivalent to the `-Wremarks` switch option)

- **#pragma diag(*warnings*)**

  This pragma can be used to restore the default behavior when neither `-w` or `-Wremarks` is specified, which is to display warnings but inhibit remarks.

Saving or Restoring the Current Behavior of All Diagnostics

This form has the following syntax:

```
#pragma diag(ACTION)
```

The effects are as follows:

- **#pragma diag(push)**

  This pragma may be used to store the current state of the severity of all diagnostic error messages.

- **#pragma diag(pop)**

  This pragma restores all diagnostic error messages that was previously saved with the most recent `push`.

All `#pragma diag(push)` directives must be matched with the same number of `#pragma diag(pop)` directives in the overall translation unit, but need not be matched within individual source files. Note that the error threshold (set by the `remarks`, `warnings` or `errors` keywords) is also saved and restored with these directives.

The duration of such modifications to diagnostic severity are from the next line following the pragma to either the end of the translation unit, the next `#pragma diag(pop)` directive, or the next overriding `#pragma diag()` directive with the same error number. These pragmas may be used anywhere and are not affected by normal scoping rules.

All command-line overrides to diagnostic severity are processed first and any subsequent `#pragma diag()` directives will take precedence, with the restore action changing the severity back to that at the start of compilation after processing the command-line switch overrides.

Note that the directives to modify specific diagnostics are singular (for example, "`error`"), and the directives to modify classes of diagnostics are plural (for example, "`errors`").

Memory Bank Pragmas

The memory bank pragmas provide additional performance characteristics for the memory areas used to hold code and data for the function.

By default, the compiler assumes that there are no external costs associated with memory accesses. This strategy allows optimal performance when the code and data are placed into high-performance internal memory. In cases where the performance characteristics of memory are known in advance, the compiler can exploit this knowledge to improve the scheduling of generated code.

Note that memory banks are different from sections:

- Section is a "hard" placement, using a name that is meaningful to the linker. If the `.LDF` file does not map the named section, a linker error occurs.
- A memory bank is a "soft" placement, using a name that is not visible to the linker. The compiler uses optimization to take advantage of the bank's performance characteristics. However, if the `.LDF` file maps the code or data to memory that performs differently, the application still functions (albeit with a possible reduction in performance).

#pragma code_bank(*bankname*)

This pragma informs the compiler that the instructions for the immediately-following function are placed in a memory bank called _bankname_. Without this pragma, the compiler assumes that the instructions are placed into a bank called "__code". When optimizing the function, the compiler takes note of attributes of memory bank *bankname*, and determines how long it takes to fetch each instruction from the memory bank.

In the example,

```
#pragma code_bank(slowmem)
int add_slowly(int x, int y) { return x + y; }

int add_quickly(int a, int b) { return a + b; }
```

the add_slowly() function is placed into the bank "slowmem", which may have different performance characteristics from the "__code" bank, into which add_quickly() is placed.

#pragma data_bank(*bankname*)

This pragma informs the compiler that the immediately-following function uses the memory bank *bankname* as the model for memory accesses for non-local data that does not otherwise specify a memory bank. Without this pragma, the compiler assumes that non-local data should use the bank "__data" for behavioral characteristics.

In the example,

```
#pragma data_bank(green)
int green_func(void)
{
      extern int arr1[32];
      extern int bank("blue") i;
      i &= 31;
      return arr1[i++];
}
int blue_func(void)
{
```

```
        extern int arr2[32];
        extern int bank("blue") i;
        i &= 31;
    return arr2[i++];
}
```

In both `green_func()` and `blue_func()`, `i` is associated with the memory bank "`blue`", and the retrieval and update of `i` are optimized to use the performance characteristics associated with memory bank "`blue`".

The array `arr1` does not have an explicit memory bank in its declaration. Therefore, it is associated with the memory bank "`green`", because `green_func()` has a specific default data bank. In contrast, `arr2` is associated with the memory bank "`__data`", because `blue_func()` does not have a `#pragma data_bank` preceding it.

#pragma stack_bank(*bankname*)

This pragma informs the compiler that all locals for the immediately-following function are to be associated with memory bank *`bankname`*, unless they explicitly identify a different memory bank. Without this pragma, all locals are assumed to be associated with the memory bank "`__stack`". In the example,

```
#pragma stack_bank(mystack)
short dotprod(int n, const short *x, const short *y)
{
       int sum = 0;
       int i = 0;
       for (i = 0; i < n; i++)
           sum += *x++ * *y++;
       return sum;
}

int fib(int n)
{
       int r;
       if (n < 2) {
          r = 1;
```

```
        } else {
           int a = fib(n-1);
           int b = fib(n-2);
           r = a + b;
        }
        return r;
}

#include <sys/exception.h>
#pragma stack_bank(sysstack)
EX_INTERRUPT_HANDLER(count_ticks)
{
        extern int ticks;
        ticks++;
}
```

The `dotprod()` function places the `sum` and `i` values into the memory bank "`mystack`", while `fib()` places `r`, `a` and `b` into the memory bank "`__stack`", because there is no `stack_bank` pragma. The `count_ticks()` function does not declare any local data, but any compiler-generated local storage make use of the "`sysstack`" memory bank's performance characteristics.

#pragma bank_memory_kind(*bankname, kind*)

This pragma informs the compiler what kind of memory the memory bank *`bankname`* is. The following kinds of memory are allowed by the compiler:

- Internal – the memory bank is high-speed in-core memory
- L2 – the memory bank is on-chip, but not in-core
- External – the memory bank is external to the processor

The pragma must appear at global scope, outside any function definitions, but need not immediately precede a function definition. In the example,

```
#pragma bank_memory_kind(blue, internal)
int sum_list(bank("blue") const int *data, int n)
```

```
{
       int sum = 0;
       while (n--)
           sum += data[n];
       return sum;
}
```

the compiler knows that all accesses to the `data[]` array are to the "`blue`" memory bank, and hence to internal, in-core memory.

#pragma bank_read_cycles(*bankname, cycles*)

This pragma tells the compiler that each read operation on the memory bank *`bankname`* requires *`cycles`* cycles before the resulting data is available. This allows the compiler to schedule sufficient code between the initiation of the read and the use of its results, to prevent unnecessary stalls.

In the example,

```
#pragma bank_read_cycles(slowmem, 20)
int dotprod(int n, const int *x, bank("slowmem") const int *y)
{
       int i, sum;
       for (i=sum=0; i < n; i++)
           sum += *x++ * *y++;
       return sum;
}
```

the compiler assumes that a read from `*x` takes a single cycle, as this is the default read time, but that a read from `*y` takes twenty cycles, because of the pragma.

The pragma must appear at global scope, outside any function definitions, but need not immediately precede a function definition.

#pragma bank_write_cycles(*bankname, cycles*)

This pragma tells the compiler that each write operation on memory bank *bankname* requires *cycles* cycles before it completes. This allows the compiler to schedule sufficient code between the initiation of the write and a subsequent read or write to the same location, to prevent unnecessary stalls.

In the following example,

```
void write_buf(int n, const char *buf)
{
      volatile bank("output") char *ptr = REG_ADDR;
      while (n--)
          *ptr = *buf++;
}
#pragma bank_write_cycles(output, 6)
```

the compiler knows that each write through `ptr` to the "`output`" memory bank takes six cycles to complete.

The pragma must appear at global scope, outside any function definitions, but need not immediately precede a function definition. This is shown in the preceding example.

#pragma bank_optimal_width(*bankname, width*)

This pragma informs the compiler that *width* is the optimal number of bits to transfer to/from memory bank *bankname* in a single cycle. This can be used to indicate to the compiler that some memories can benefit from vectorization and similar strategies more than others. The *width* parameter must be 8, 16, 24 or 32.

In the example,

```
void memcpy_simple(char *dst, const char *src, size_t n)
{
      while (n--)
          *dst++ = *src++;
```

```
}
#pragma bank_optimal_width(__code, 16)
```

the compiler knows that the instructions for the generated function would be best fetched in multiples of 16 bits, and so can select instructions accordingly.

The pragma must appear at global scope, outside any function definitions, but need not immediately precede a function definition. This is shown in the preceding example.

GCC Compatibility Extensions

The compiler provides compatibility with the C dialect accepted by version 3.2 of the GNU C Compiler. Many of these features are available in the C99 ANSI Standard. A brief description of the extensions is included in this section. For more information, refer to the following web address:

`http://gcc.gnu.org/onlinedocs/gcc-3.2.1/gcc/C-Extensions.html#C%20Extensions`

The GCC compatibility extensions are only available in C dialect mode. They are not accepted in C++ dialect mode.

Statement Expressions

A statement expression is a compound statement enclosed in parentheses. A compound statement itself is enclosed in braces { }, so this construct is enclosed in parentheses-brace pairs ({ }).

The value computed by a statement expression is the value of the last statement (which should be an expression statement). The statement expression may be used where expressions of its result type may be used. But they are not allowed in constant expressions.

Statement expressions are useful in the definition of macros as they allow the declaration of variables local to the macro. In the following example,

```
#define min(a,b) ({                 \
   short __x=(a),__y=(b),__res;     \
   if (__x > __y)                   \
      __res = __y;                  \
   else                             \
   __res = __x;                     \
   __res;                           \
})

int use_min() {
   return  min(foo(), thing()) + 2;
}
```

The `foo()` and `thing()` statements get called once each because they are assigned to the variables `__x` and `__y` which are local to the statement expression that `min` expands to. The `min()` can be used freely within a larger expression because it expands to an expression.

Labels local to a statement expression can be declared with the `__label__` keyword. For example,

```
({
   __label__ exit;
   int i;
   for (i=0; p[i]; ++i) {
      int d = get(p[i]);
      if (!check(d)) goto exit;
      process(d);
   }
exit:
   tot;
})
```

Statement expressions are not supported in C++ mode. Statement expressions are an extension to C originally implemented in the GCC compiler. Analog Devices support the extension primarily to aid porting code written for that compiler.

When writing new code, consider using inline functions, which are compatible with ANSI/ISO standard C++ and C99, and are as efficient as macros when optimization is enabled.

Type Reference Support Keyword (Typeof)

The `typeof( expression )` construct can be used as a name for the type of expression without actually knowing what that type is. It is useful for making source code that is interpreted more than once, such as macros or include files, more generic. The `typeof` keyword may be used where ever a `typedef` name is permitted such as in declarations and in casts.

For example,

```
#define abs(a) ({                   \
   typeof(a) __a = a;               \
   if (__a < 0) __a = - __a;        \
   __a;                             \
})
```

shows `typeof` used in conjunction with a statement expression to define a "generic" macro with a local variable declaration.

The argument to `typeof` may also be a type name. Because `typeof` itself is a type name, it may be used in another `typeof( type-name )` construct. This can be used to restructure the C-type declaration syntax. For example,

```
#define pointer(T)   typeof(T *)
#define array(T, N)  typeof(T [N])

array (pointer (char), 4) y;
```

declares `y` to be an array of four pointers to `char`.

The `typeof` keyword is not supported in C++ mode.
The `typeof` keyword is an extension to C originally implemented in the GCC compiler. It should be used with caution because it is not compatible with other dialects of C or C++ and has not been adopted by the more recent C99 standard.

GCC Generalized Lvalues

A cast is an `lvalue` (may appear on the left-hand side of an assignment) if its operand is an `lvalue`. This is an extension to C, provided for compatibility with GCC. It is not allowed in C++ mode.

A comma operator is an `lvalue` if its right operand is an `lvalue`. This is an extension to C, provided for compatibility with GCC. It is a standard feature of C++.

A conditional operator is an `lvalue` if its last two operands are `lvalues` of the same type. This is an extension to C, provided for compatibility with GCC. It is a standard feature of C++.

Conditional Expressions With Missing Operands

The middle operand of a conditional operator can be left out. If the condition is nonzero (true), then the condition itself is the result of the expression. This can be used for testing and substituting a different value when a pointer is NULL. The condition is only evaluated once; therefore, repeated side effects can be avoided. For example,

```
printf("name = %s\n",  lookup(key)?:"-");
```

calls `lookup()` once, and substitutes the string "-" if it returns NULL. This is an extension to C, provided for compatibility with GCC. It is not allowed in C++ mode.

Hexadecimal Floating-Point Numbers

C99 style hexadecimal floating-point constants are accepted. They have the following syntax.

```
hexadecimal-floating-constant:
  {0x|0X} hex-significand binary-exponent-part [ floating-suffix ]
hex-significand: hex-digits [ . [ hex-digits ]]
binary-exponent-part: {p|P} [+|-] decimal-digits
floating-suffix: { f | l | F | L }
```

The hex-significand is interpreted as a hexadecimal rational number. The digit sequence in the exponent part is interpreted as a decimal integer. The exponent indicates the power of two by which the significand is to be scaled. The floating suffix has the same meaning that it has for decimal floating constants—a constant with no suffix is of type `double`, a constant with suffix `F` is of type `float`, and a constant with suffix `L` is of type `long double`.

Hexadecimal floating constants enable the programmer to specify the exact bit pattern required for a floating-point constant. For example, the declaration

```
float f = 0x1p-126f;
```

causes `f` to be initialized with the value `0x800000`.

Hexadecimal floating constants are not supported in C++ mode.

Zero-Length Arrays

Arrays may be declared with zero length. This is an anachronism supported to provide compatibility with GCC. Use variable-length array members instead.

Variable Argument Macros

The final parameter in a macro declaration may be followed by dots (...) to indicate the parameter stands for a variable number of arguments.

For example,

```
#define trace(msg, args...)   fprintf (stderr, msg, ## args);
```

can be used with differing numbers of arguments,

```
trace("got here\n");
trace("i = %d\n", i);
trace("x = %f, y = %f\n", x, y);
```

The *##* operator has a special meaning when used in a macro definition before the parameter that expands the variable number of arguments: if the parameter expands to nothing, then it removes the preceding comma.

The variable argument macro syntax comes from GCC. It is not compatible with C99 variable argument macros and is not supported in C++ mode.

Line Breaks in String Literals

String literals may span many lines. The line breaks do not need to be escaped in any way. They are replaced by the character `\n` in the generated string. This extension is not supported in C++ mode. The extension is not compatible with many dialects of C, including ANSI/ISO C89 and C99. However, it is useful in `asm` statements, which are intrinsically non-portable.

Arithmetic on Pointers to Void and Pointers to Functions

Addition and subtraction is allowed on pointers to `void` and pointers to functions. The result is as if the operands had been cast to pointers to `char`. The `sizeof` operator returns one for `void` and function types.

Cast to Union

A type cast can be used to create a value of a union type, by casting a value of one of the union member's types.

Ranges in Case Labels

A consecutive range of values can be specified in a single case by separating the first and last values of the range with ....

For example,

```
case 200 ... 300:
```

Declarations Mixed With Code

In C mode, the compiler accepts declarations placed in the middle of code. This allows the declaration of local variables to be placed at the point where they are required. Therefore, the declaration can be combined with initialization of the variable.

For example, in the following function

```
void func(Key k) {
Node *p = list;
while (p && p->key != k)
   p = p->next;
if (!p)
   return;
Data *d = p->data;
while (*d)
   process(*d++);
}
```

the declaration of `d` is delayed until its initial value is available, so that no variable is uninitialized at any point in the function.

Escape Character Constant

The character escape "`\e`" may be used in character and string literals and maps to the ASCII Escape code, 27.

Alignment Inquiry Keyword (__alignof__)

The `__alignof__ (type-name)` construct evaluates to the alignment required for an object of a type. The `__alignof__` *`expression`* construct can also be used to give the alignment required for an object of the *`expression`* type.

If expression is an `lvalue` (may appear on the left-hand side of an assignment), the alignment returned takes into account alignment requested by pragmas and the default variable allocation rules.

(asm) Keyword for Specifying Names in Generated Assembler

The `asm` keyword can be used to direct the compiler to use a different name for a global variable or function (see also "#pragma linkage_name identifier" on page 1-161). For example,

```
int N asm("C11045");
```

instructs the compiler to use the label C11045 in the assembly code it generates wherever it needs to access the source level variable `N`. By default, the compiler would use the label `_N`.

The `asm` keyword can also be used in function declarations but not function definitions. However, a definition preceded by a declaration has the desired effect. For example,

```
extern int f(int, int) asm("func");

int f(int a, int b) {
. . .
}
```

Function, Variable and Type Attribute Keyword (__attribute__)

The __attribute__ keyword can be used to specify attributes of functions, variables and types, as in these examples,

```
void func(void) __attribute__ ((section("fred")));
int a __attribute__ ((aligned (8)));
typedef struct {int a[4];} __attribute__((aligned (4))) Q;
```

Since the __attribute__ keyword is supported, the code, written for GCC, can be ported. All attributes accepted by GCC on ix86 are accepted. The ones that are actually interpreted by the compiler are described in the sections of this manual describing the corresponding pragmas (see "Pragmas" on page 1-148).

Preprocessor-Generated Warnings

The preprocessor directive #warning causes the preprocessor to generate a warning and continue preprocessing. The text that follows the #warning directive on the line is used as the warning message. For example,

```
#ifndef __ADSPBLACKFIN__
#warning   This program is written for Blackfin processors
#endif
```

Blackfin Processor-Specific Functionality

This section provides information about functionality that is specific to the Blackfin processors.

This section describes:

- "Startup Code Overview"
- "Support for argv/argc" on page 1-197
- "Profiling With Instrumented Code" on page 1-197
- "Controlling Available Memory Size" on page 1-201
- "Interrupt Handler Support" on page 1-201
- "Caching and Memory Protection" on page 1-208

Startup Code Overview

Startup code, which is invoked when the processor starts running, initializes a default environment before calling `main()`. The VisualDSP++ 4.0 **Project Wizard** can be used to generate startup code based on the configuration options selected.

Alternatively, default startup code is contained in the run-time libraries, the source for which is in the `...Blackfin\lib\src\libc\crt\basiccrt.s` file. You may want to modify this code to suit your specific target environment, or perhaps even replace it completely.

Refer to "C/C++ Run-Time Header and Startup Code" on page 1-232 for more information.

Support for argv/argc

By default, the facility to specify arguments that get passed to your `main()` (`argv/argc`) at runtime is enabled. However, to correctly set up `argc` and `argv` requires additional configuration by the user. You need to modify your application in the following ways:

1. Define your command-line arguments in C by defining a variable called "`__argv_string`". When linked, your new definition overrides the default zero definition otherwise found in the C run-time library. For example,

```
const char __argv_string[] = "-in x.gif -out y.jpeg";
```

2. To use command-line arguments as part of Profile-Guided Optimizations (PGO), it is necessary to define `__argv_string` within a memory section called `MEM_ARGV`. Therefore, define a memory section called `MEM_ARGV` in your `.LDF` file and include the definition of `___argv_string` in it if you are using PGO. The default `.LDF` files do this for you if macro `IDDE_ARGS` is defined at link time.

Profiling With Instrumented Code

The profiling facilities determine how many times each function is called and how many cycles are used while the function is active. The information is gathered by an additional library linked into the executable file. The profiling routine is invoked by additional function calls at the start and end of each function. The compiler inserts these extra calls when profiling is enabled.

The compiler profiling facilities should not be confused with similar functionality in the simulator, which works on a per-instruction basis rather than on a per-function basis.
The compiler profiling facilities are designed for single-threaded systems, and do not work if function invocations from more than one thread are in progress concurrently.

Generating Instrumented Code

The `-p[1|2]` switch (on page 1-47) turns on the compiler's profiling facility when converting C/C++ source into assembly code. The compiler cannot instrument assembly files or files that have already been compiled to object files.

- The `-p1` option writes accumulated profile data to the file "`mon.out`"in the current directory.
- The `-p2` option writes accumulated profile data to standard output.
- The `-p` option writes accumulated profile data to both standard output and the `mon.out` file in the current directory.

Running the Executable

The executable may produce two forms of output. The first is a dump of data to standard output once the program completes (generated by `-p` and `-p2`). This output lists the approximate address of each profiled function, how many times the function was invoked, and the inclusive and exclusive cycle counts.

- Exclusive cycle counts include only the cycles spent processing the function.
- Inclusive cycle counts also include the sum total of cycle counts in any function invoked from this specified function.

For example, in the following program

```
int apples, bananas;
void apple(void) {
     apples++;       // 10 cycles
}

void banana(void) {
     bananas++;      // 10 cycles
```

```
     apple();         // 10 cycles
}                     // 20 cycles total

int main(void) {
    apple();          // 10 cycles
    apple();          // 10 cycles
    banana();         // 20 cycles
    return 0;         // 40 inclusive cycles total
}                     // + exclusive cycles for main itself
```

assume that `apple()` takes 10 cycles per call and assume that `banana()` takes 20 cycles per call, of which 10 are accounted for by its call to `apple()`. The program, when run, calls `apple()` three times: twice directly and once indirectly through `banana()`. The `apple()` function clocks up 30 cycles of execution, and this is reported for both its inclusive and exclusive times, since `apple()` does not call other functions.

The `banana()` function is called only once. It reports 10 cycles for its exclusive time, and 20 cycles for its inclusive time. The exclusive cycles are for the time when `banana()` is incrementing `bananas` and is not "waiting" for another function to return, and so it reports 10 cycles. The inclusive cycles include these 10 exclusive cycles and also include the 10 cycles `apple()` used when called from `banana()`, giving a total of 20 inclusive cycles.

The `main()` function is called only once, and calls three other functions (`apple()` twice, `banana()` once). Between them, `apple()` and `banana()` use up to 40 cycles, which appear in the `main()`'s inclusive cycles. The `main()`'s exclusive cycles are for the time when `main()` is running, but is not in the middle of a call to either `apple()` or `banana()`.

The second form of output is a file in the current directory called `mon.out` (`-p` and `-p1`). The `mon.out` is a binary file that contains a copy of the data written to standard output. There is no way to change the file name used.

Post-Processing mon.out File

The `profblkfn.exe` program is a Windows program that processes the contents of the `mon.out` file. It reads both the `mon.out` file and the `.DXE` file that had produced it. It displays the cycle counts along with the names of the functions recorded in the `mon.out` file associated with the counts. The `profblkfn` program is invoked as:

```
profblkfn prog.dxe
```

Specify the `.DXE` file only. The `mon.out` file must be present in the current directory and must be produced by the named `.DXE` file.

Computing Cycle Counts

When profiling is enabled, the compiler instruments the generated code by inserting calls to a profiling library at the start and end of each compiled function. The profiling library samples the processor's cycle counter and records this figure against the function just started or just completed.

The profiling library itself consumes some cycles, and these overheads are not included in the figures reported for each function, so the total cycles reported for the application by the profiler will be less than the cycles consumed during the life of the application. In addition to this overhead, there is some approximation involved in sampling the cycle counter, because the profiler cannot guarantee how many cycles will pass between a function's first instruction and the sample. This is affected by the optimization levels, the state preserved by the function, and the contents of the processor's pipeline. The profiling library knows how long the call entry and exit takes "on average", and adjusts its counts accordingly.

Because of this adjustment, profiling using instrumented code provides an approximate figure, with a small margin for error. This margin is more significant for functions with a small number of instructions than for functions with a large number of instructions.

Controlling Available Memory Size

The heap size is specified in the .LDF file in the VisualDSP\Blackfin\LDF directory. By default, the compiler uses the file arch.ldf, where *arch* is specified via the -proc *arch* switch. For example, if -proc ADSP-BF533 is used, the compiler defaults to using adsp-BF533.ldf. The entry controlling the heap has a format similar to

```
MEM_HEAP { TYPE(RAM) START(0xFF804000) END(0xFF807DFF) WIDTH(8) }
```

The actual values specified in the .LDF file should reflect the memory map available on the actual system.

Internally, malloc() uses the _Sbrk() library function to obtain additional space from the HEAP system. The start and end addresses of the HEAP segment can be changed to give a larger or smaller heap, and the library will adjust accordingly. If the segment size is increased, the surrounding segments must be decreased accordingly; otherwise, memory corruption may occur. See "Using Multiple Heaps" on page 1-245 for more information.

Interrupt Handler Support

The Blackfin C/C++ compiler provides support for interrupts and other events used by the Blackfin processor architecture (see Table 1-14).

The Blackfin system has several different classes of events, not all of which are supported by the ccblkfn compiler. Handlers for these events are called Interrupt Service Routines (ISRs).

Table 1-14. System Events

Event	Priority	Supported
Emulation	Highest	No
Reset		Yes
NMI		Yes
Exception		Yes
Interrupts	Lowest	Yes

Resets are supported by treating them like general-purpose interrupts for code-generation purposes. This means that the C/C++ compiler supports interrupt, exception, and NMI events.

The compiler provides facilities for defining an `ISR` function, registering it as an event handler, and for obtaining the saved processor context.

Defining an ISR

To define a function as an ISR, the `sys/exception.h` header must be included and the function must be declared and defined using macros defined within this header file. There is a macro for each of the three kinds of events the compiler supports:

```
EX_INTERRUPT_HANDLER
EX_EXCEPTION_HANDLER
EX_NMI_HANDLER
```

By default, the ISRs generated by the compiler are not re-entrant; they disable the interrupt system on entry, and re-enable it on exit. You may also define ISRs for interrupts which are re-entrant, and which re-enable the interrupt system soon after entering the ISR.

There is a different macro for specifying a re-entrant interrupt handler:

```
EX_REENTRANT_HANDLER
```

For example, the following code

```
#include <sys/exception.h>
static int number_of_interrupts;

EX_INTERRUPT_HANDLER(my_isr)
{
      number_of_interrupts++;
}
```

declares and defines `my_isr()` as a handler for interrupt-type events (for example, the routine returns using an `RTI` instruction).

The macro used for defining the ISR is also suitable for declaring it as a prototype:

```
EX_INTERRUPT_HANDLER(my_isr);
```

The `EX_INTERRUPT_HANDLER()` macro uses a generic pragma, `#pragma interrupt`, to indicate that the function is an interrupt handler. This generic pragma does not indicate which interrupt the function handles. The `-workaround isr-imask-check` switch selection (on page 1-62) for the hardware anomaly 05-00-0071 on the ADSP-BF535 processor requires explicit information on the level of interrupt being handled, so that the interrupt can be re-raised if the interrupt is taken while a `CLI` instruction is being committed.

Such an ISR is defined as:

```
#pragma interrupt_level_6
void my_handler(int _r0x, int _r1x, int _r2x) {
   int_count++;
}
```

There are eleven such level-specific pragmas, from 5 through to 15, corresponding to the Blackfin event table entries for interrupts.

If the `isr-imask-check` workaround is enabled, ISRs declared without explicit interrupt levels—such as those declared using `EX_INTERRUPT_HANDLER()`—check for interrupts occurring while a CLI is committed and return immediately if this is detected. They do not attempt to re-raise the interrupt.

While thread-safe variants of the C/C++ run-time libraries exist, many functions are not interrupt-safe as they access global data structures. It is therefore recommended that ISRs do not make library function calls, as unexpected behavior may result if the interrupt occurs during a call to such a function. An alternative approach is to disable interrupts before the application makes run-time library calls. This may be disadvantageous for

time-critical applications as interrupts may be disabled for a long period of time. The DSP run-time library functions do not modify global data structures and are therefore interrupt-safe.

Registering an ISR

ISRs, once defined, can be registered in the Event Vector Table (EVT) using the `register_handler()` function. This function operates in a similar manner to the UNIX `signal()` function.

It takes two parameters, defining the event and the ISR, and returns the previously registered ISR (if any). The event is specified using the `interrupt_kind` enumeration from `exception.h`. For example,

```
typedef enum {
   ik_emulation, ik_reset, ik_nmi, ik_exception,
   ik_global_int_enable, ik_hardware_err, ik_timer,ik_ivg7,
   ik_ivg8, ik_ivg9, ik_ivg10, ik_ivg11, ik_ivg12, ik_ivg13,
   ik_ivg14, ik_ivg15
} interrupt_kind;
ex_handler_fn register_handler(interrupt_kind, ex_handler_fn);
```

There are two special values that can be passed to `register_handler()` in place of real ISRs:

- `EX_INT_IGNORE` installs a handler that "ignores" the event and immediately returns from the event.
- `EX_INT_DEFAULT` installs the default handler. The default handler invokes the currently registered handler for the corresponding ANSI C signal, as described in "ISRs and ANSI C Signals".

ISRs and ANSI C Signals

ISRs provide similar functionality to ANSI C signal handlers, and their behavior is related. An ISR is a function that can be registered directly in the processor's Event Vector Table (EVT). The `ISR` function saves its own context, as required. In contrast, an ANSI C signal handler is a normal C

function that has been registered as a handler. When an event occurs, some other dispatcher must save the processor context before invoking the signal handler.

ISRs and signal handlers are not interchangeable. A signal handler cannot act as an ISR, because it does not save or restore the context, nor it terminates with the correct return instruction. An ISR cannot act as a signal handler, because it terminates the event directly rather than returning to the dispatcher.

When a signal handler is installed, a default ISR is also installed in the EVT which invokes the signal handler when the event occurs. When the `raise()` function is used to invoke a signal handler explicitly, `raise()` actually generates the corresponding event (if possible). This causes the ISR to invoke the signal handler.

You may choose to install normal C functions as signal handlers, or to register ISRs directly, but you should not do both for a given event.

The ANSI C signals are registered using `signal()`, or using the Analog Devices extension `interrupt()`, unlike ISRs, which are registered using `register_handler()`.

Saved Processor Context

When generating code for an ISR, the compiler creates a prologue that saves the processor context on the supervisor stack. This context is accessible to the ISR. The `exception.h` file defines a structure, `interrupt_info`, that contains fields for all the information that defines the kind of event that occurred and for the values of all the registers that were saved during the prologue.

For a list of saved registers, see "Fetching Saved Registers" on page 1-207.

There are two facilities for gaining access to the event context:

- `get_interrupt_info()` function
- `SAVE_REGS()` macro (see "Fetching Saved Registers")

Fetching Event Details

The following function fetches the information about the event that occurred:

```
void get_interrupt_info(interrupt_kind, interrupt_info *)
```

The sort of data retrieved includes the value of `EXCAUSE` and addresses that cause exceptions faults. Note that at present, the function needs to be told which kind of event it is investigating.

The structure contains:

```
interrupt_kind kind;
int value;
void *pc;
void *addr;
unsigned status;
interrupt_regs regs;
interrupt_regs *regsaddr;
```

These fields are set as:

- **Exceptions**
  The `pc` is set to the value of `RETX`, and `value` is set to the value of `SEQSTAT`.

  For exceptions that involve address faults, address and status are set to the values of the Memory Mapping Registers (MMRs) for `DATA_FAULT_ADDR` and `DATA_FAULT_STATUS` or `CODE_FAULT_ADDR` and `CODE_FAULT_STATUS`, as appropriate.

- **Hardware Errors**
  The `pc` is set to the value of `RETI`, and `value` is set to the value of `SEQSTAT`.
- **NMI Events**
  The `pc` is set to the value of `RETN`.
- **All Other Events**
  The `pc` is set to the value of `RETI`.

Fetching Saved Registers

The following macro obtains a copy of the registers saved during the ISR prologue:

```
SAVE_REGS(interrupt_info *)
```

It also sets `regsaddr` in the `interrupt_info` structure to point to the start of the saved registers on the supervisor stack. Therefore, any changes made using `regsaddr` within the ISR are reflected in the processor state when it is restored by the ISR epilogue.

The following registers are always saved during the ISR prologue. They are accessible through the saved context.

- All `DREGS` (`R0`, `R1`, `R2`, `R3`, `R4`, `R5`, `R6`, `R7`)
- All `PREGS` (`P0`, `P1`, `P2`, `P3`, `P4`, `P5`)
- Frame pointer (`FP`)
- Arithmetic status (`ASTAT`)

Additional registers are saved as required, depending on the resources used by the ISR. These registers are not accessible through the saved context.

The registers that are optionally saved include:

- Hardware loop registers (`LB0, LB1, LT0, LT1, LC0, LC1`)
- Accumulators (`A0w, A1w, A0x, A1x`)
- Circular buffer registers (`I0-3, L0-3, B0-3, M0-3`)

When using the interrupt function on a ADSP-BF533 processor, it takes approximately 100 cycles from the first instruction of the interrupt despatcher to the first instruction of user-provided code in the interrupt routine. It takes approximately 59 cycles to return from the last instruction of user-provided code in the interrupt routine to the original path of execution.

For the ADSP-BF535 processor, it takes approximately 149 cycles from the first instruction of the despatcher to the first user-provided instruction of the interrupt routine. It takes approximately 228 cycles to return from the last instruction of user-provided code in the interrupt routine to the original path of execution.

The cycle count figures may vary depending on the registers that are required to be preserved for the execution of the interrupt servicing routine.

Caching and Memory Protection

The Blackfin processors support the caching of external memory or L2 SRAM (where available) into L1 SRAM, for both instruction and data memory. Caching can eliminate much of the performance penalty of using external memory with minimal effort on the application developer's part.

The Blackfin processor caches are configurable devices. Instruction and data caches can be enabled together or separately, and the memory spaces they cache are configured separately. The cache configuration is defined through the memory protection hardware, using tables that define Cache Protection Lookaside Buffers (CPLBs). These CPLBs define the start

addresses, sizes, and attributes of areas of memory for which memory accesses are permitted (including whether the area of memory is to be cached).

Refer to the *Hardware Reference* of the appropriate Blackfin processor for specific details.

The Blackfin run-time library provides support for cache configuration by providing routines that can be used to initialize and maintain the CPLBs from a configuration table.

Both the **Project Wizard**-generated C/C++ run-time headers (CRTs) and default pre-compiled CRT objects make use of these library routines. The default configuration does not enable CPLBs. The support routines are designed such that they can easily be incorporated into users' systems, and so that the configuration can be turned on or off via a debugger, without the need for re-linking the application. (See the *VisualDSP++ 4.0 User's Guide* and "C/C++ Run-Time Header and Startup Code" on page 1-232 for more information.)

___cplb_ctrl Control Variable

The CPLB support is controlled through a global integer variable, `___cplb_ctrl`. Its C name has two leading underscores and its assembler name has three leading underscores. The value of this variable determines whether the startup code enables the CPLB system. By default, the variable has the value "zero", indicating that CPLBs should not be enabled.

The variable's value is a bitmask, based on the macros defined in the `<cplb.h>` header. The macros are:

- `CPLB_ENABLE_ICPLBS` – turns on instruction CPLBs
- `CPLB_ENABLE_ICACHE` – turns on instruction caching into L1 Instruction memory
- `CPLB_ENABLE_DCPLBS` – turns on data CPLBs

- CPLB_ENABLE_DCACHE – turns on data caching into L1 Data A memory
- CPLB_ENABLE_DCACHE2 – turns on data caching into L1 Data B memory
- CPLB_SET_DCBS – sets the Data Cache Bank Select bit in the DMEM_CONTROL register. This specifies which bit of a memory address determines the data cache bank (A or B) used to cache the location. Depending on the placement of data within the application memory space, one setting or the other ensures more data is cached at runtime. This bit has no effect unless both CPLB_ENABLE_DCACHE and CPLB_ENABLE_DCACHE2 bits are also set. See the hardware reference manual for further details.

These macros are OR'ed together to produce the value for ___cplb_ctrl. Note that

- If CPLB_ENABLE_DCACHE2 is set, CPLB_ENABLE_DCACHE must also be set.
- If any of the three cache bits are set, the corresponding CPLB_ENABLE_ICPLBS or CPLB_ENABLE_DCPLBS bit must also be set.

There is a default definition of ___cplb_ctrl in the C run-time library, which defaults to disabling CPLBs and caching. This default definition is overridden by any definition in the CRT startup code generated by the **Project Wizard**, or alternatively by providing your own definition within your application. For example,

```
#include <stdio.h>
#include <cplb.h>
int __cplb_ctrl =         // C syntax with two underscores
   CPLB_ENABLE_ICPLBS |
   CPLB_ENABLE_ICACHE;
int main(void) {
   printf("Hello world\n");
   return 0;
}
```

The new definition enables CPLBs and turns on instruction caching; data caching is not enabled.

CPLB Installation

When ___cplb_ctrl indicates that CPLBs are to be enabled, the startup code calls the routine _cplb_init. This routine sets up instruction and data CPLBs from a table, and enables the memory protection hardware.

There are sixteen CPLBs for each instruction and data space. On a simple system, this is sufficient, and _cplb_init installs all available CPLBs from its configuration table into the active table. On more complex systems, there may need to be more CPLBs than can be active at once. In such systems, a time may come when the application attempts to access memory that is not covered by one of the active CPLBs. This raises a CPLB miss exception.

For these occasions, the library includes a CPLB management routine _cplb_mgr. This routine should be called from an exception handler that has determined a CPLB miss has occurred (a data miss or an instruction miss). The _cplb_mgr routine identifies which inactive CPLB needs to be installed to resolve the access, and replaces one of the active CPLBs with this one.

If CPLBs are to be enabled, the default startup code installs a default exception handler called _cplb_hdr; this does nothing except test for CPLB miss exceptions, which it delegates to _cplb_mgr. It is expected that users have their own exception handlers that deal with additional events.

Since the CPLB configuration tables and management code need to be present during all CPLB miss exceptions, these are placed into a separate "cplb_code" and "cplb_data" sections, and the CPLBs that refer to these sections must be:

- flagged as being "locked", so they are not replaced by inactive CPLBs during misses

- among the first sixteen CPLBs, so they are loaded into the active table during initialization

Cache Configurations

Although CPLBs may be used for their protection capabilities, they are most commonly used to enable caching. The `___cplb_ctrl` variable is the means by which the application directs the run-time library to install CPLBs for caching. The library defines the following configurations, although not all configurations may be available on all Blackfin processors:

- No cache
- L1 SRAM Instruction as cache
- L1 SRAM Data A as cache
- L1 SRAM Data A and B as cache
- L1 SRAM Instruction and Data A as cache
- L1 SRAM Instruction, Data A and Data B as cache

Note that if any cache is enabled, the corresponding data or instruction CPLBs must also be enabled. Furthermore, if you are using the default `.LDF` files, you must also tell the linker that the cache is enabled; this is discussed in more detail in "Default Cache Configuration" and "LDFs and Cache" on page 1-216.

If any cache is enabled, the respective caches are set up during `_cplb_init`, using the CPLB configuration tables. On the ADSP-BF535 and AD6532 processors, if the cache is enabled, the current cache contents are invalidated using the functions described in "Cache Invalidation" on page 1-214. With other Blackfin processors, the cache is automatically invalidated at power-up.

Default Cache Configuration

Although the default value for ___cplb_ctrl is that no cache or CPLBs are enabled, the default system contains CPLB configuration tables that permit caching. The default configuration tables supplied differ for the parts available.

The default configuration tables are defined in files called cplbtab*n*.s in VisualDSP\Blackfin\lib\src\libc\crt, where *n* is the part number.

Blackfin Processor	Configuration file
ADSP-BF531	cplbtab531.s
ADSP-BF532	cplbtab532.s
ADSP-BF533	cplbtab533.s
ADSP-BF534	cplbtab534.s
ADSP-BF535	cplbtab535.s
ADSP-BF536	cplbtab536.s
ADSP-BF535	cplbtab535.s
ADSP-BF538	cplbtab538.s
ADSP-BF539	cplbtab539.s
ADSP-BF561 (Core A)	cplbtab561a.s
ADSP-BF561 (Core B)	cplbtab561b.s
ADSP-BF566 (Core A)	cplbtab566a.s
ADSP-BF566 (Core B)	cplbtab566b.s
AD6532	cplbtab6532.s

Each file defines two tables:

1. icplbs_table[] – the Instruction CPLBs
2. dcplbs_table[] – the Data CPLBS

The table's structure is defined by `<cplbtab.h>`, specifying the start address of each area of memory, and the controlling attributes for that area.

The default tables include areas of memory for L1 SRAM, Internal L2 (where present), external Asynchronous and SDRAM memory, and other memory spaces. The external areas are configured to be cacheable using the Write Through mode by default. If no cache is enabled, but CPLBs are enabled, the run-time library masks off the cacheable flags on the CPLBs before making them active.

Changing Cache Configuration

The value of `___cplb_ctrl` may be changed in several ways:

- The **Project Wizard** can be used to generate CRT startup code that includes a definition of the `___cplb_ctrl` variable based on the selected cache configuration.
- It may be defined as a new global variable with an initialization value. This definition supersedes the definition in the library. The example given in "___cplb_ctrl Control Variable" on page 1-209 uses this approach.
- The linked-in version of the variable may be altered in a debugger, after loading the application but before running it, so that the startup code sees a different value.

Cache Invalidation

The `cache_invalidate` routine may be used to invalidate the processor's instruction and/or data caches. It is defined as:

```
#include <cplbtab.h>
void cache_invalidate(int cachemask);
```

Its parameter is a bitmask indicating which caches should be cleared.

Bit set	Cache invalidated
CPLB_ENABLE_ICACHE	Instruction Cache
CPLB_ENABLE_DCACHE	Data Cache A
CPLB_ENABLE_DCACHE2	Data Cache B

The cache_invalidate routine uses several supporting routines:

```
#include <cplbtab.h>
void icache_invalidate(void);
void dcache_invalidate(int a_or_b);
void dcache_invalidate_both(void);
```

The icache_invalidate routine clears the instruction cache.

The dcache_invalidate routine clears a single data cache, determined by the a_or_b parameter.

CPLB_INVALIDATE_A invalidates Data Cache A
CPLB_INVALIDATE_B invalidates Data Cache B

The dcache_invalidate_both routine clears both Data Cache A and Data Cache B. On the ADSP-BF535 and AD6532 processors, it does this by calling dcache_invalidate for each cache. On other Blackfin processors, it toggles control bits in the DMEM_CONTROL register, which invalidates the contents of both data caches in a single operation.

The dcache_invalidate and dcache_invalidate_both routines do not flush any modified cache entries to memory first, if any memory pages are cached in Write Back mode. To flush such data prior to invalidation, use the functions described in "Cache Flushing" on page 1-220.

LDFs and Cache

The default .LDF files supplied with VisualDSP++ are designed to support caching with minimal effort.

The default LDFs have three basic configurations:

1. No external SDRAM, and no caching. All code and data are placed into internal SRAM. This is the default configuration.
2. External SDRAM (32MB), no caching. Code and data is placed both into internal SRAM and external SDRAM (32MB). Code and data is placed into internal SRAM where possible. This configuration is enabled by passing the -MDUSE_SDRAM flag to the linker at link time.
3. External SDRAM (32MB), and code and data caching. Some code and data is placed into external SDRAM Bank 0. This configuration is enabled by passing the -MDUSE_CACHE flag to the linker at link time.

The second configuration is not discussed here. The various Blackfin processors differ in their available SRAM resources, so the USE_CACHE configuration either allocates both Data Bank A and Data Bank B for cache, or just Data Bank A.

When enabling caches using ___cplb_ctrl, it is imperative that USE_CACHE also be specified. Consider the following scenarios:

___cplb_ctrl settings	USE_CACHE settings	Result
No Cache	No Cache	**Higher performance:** Code/data mapped into L1 SRAM and used correctly.
Use Cache	Use Cache	**Higher performance:** Code/data not mapped into L1 SRAM, leaving L1 SRAM free for caching.

___cplb_ctrl settings	USE_CACHE settings	Result
No Cache	Use Cache	**Lower performance:** Code/data not mapped into L1 SRAM, but L1 SRAM not used for cache. Wasted resources and poor performance.
Use Cache	No Cache	**Lower performance:** Code/data mapped into L1 SRAM, and L1 SRAM used as cache: code/data is corrupted by caching process.

Therefore, if the `.LDF` file believes cache is to be used, but `___cplb_ctrl` specifies otherwise, resources are wasted, but the application still functions. In fact, this is the case if only code or data caches are requested by `___cplb_ctrl`, but not both.

The last entry of the above table shows a configuration that must be avoided since mapping code/data into L1 SRAM, which is then configured as cache, leads to corrupt code/data. This scenario can also be difficult to debug, so the run-time library provides a mechanism for protecting against this case, such as:

- The default `.LDF` files define global "guard" symbols, setting their addresses to be either 0 or 1, according to whether `USE_CACHE` is defined at link time. If objects are mapped into a cache area during linking, the guard symbol is set to 0 (indicating this cache area is not available), otherwise it is set to 1 (indicating that the cache area is available).
- When `_cplb_init` is enabling CPLBs and cache, it tests the guard symbols. If a cache has been requested via `___cplb_ctrl`, but the corresponding guard symbol indicates that the cache area has already been allocated during linktime, the library signals an error. It does so by jumping to an infinite loop of the form:

```
verbose_label_expressing_error_condition:
   IDLE;
```

```
do_not_know_what_to_do:
   JUMP verbose_label_expressing_error_condition;
```

The guard symbols have the following names:

```
___l1_code_cache
___l1_data_cache_a
___l1_data_cache_b
```

CPLB Replacement and Cache Modes

As previously noted, there may be no more than sixteen CPLBs active concurrently for each instruction space or data space. Large applications may need to address more memory than this, and may eventually access a memory location not covered by the currently-active CPLBs. At this point, a CPLB "miss exception" occurs, and the application's exception handler must select one of the active CPLBs for removal to make way for a new CPLB that covers the address being accessed. This victimization and replacement process is handled by the `_cplb_mgr` routine within the run-time library. The process varies depending on what cache modes are active.

Blackfin processors support two variants of caching: Write Through and Write Back modes.

- In Write Through mode, writes to cached memory are written both to the cache and to the memory location. Consequently, Write Through mode primarily provides performance gains for memory reads. The memory location is kept up-to-date.
- In Write Back mode, writes to cached memory are only written to the cache. They are not written to the memory location until the cache line is victimized (by an access to another memory location) or flushed (through programmatic means). The memory location is not always as up-to-date as the cache copy, and if Write Back CPLBs are victimized, there are additional overheads for flushing.

Consequently, Write Back mode provides performance gains for both reads and writes, but at the cost of increased complexity if CPLBs are victimized.

The cache mode in question—OFF, Write Through or Write Back—is specified on a per-CPLB basis, so one page may be cached in Write Through mode, another in Write Back mode, and a third not cached at all. When pages are not cached, or cached in Write Through mode, the victimization and replacement process is relatively simple. When pages are cached in Write Back mode, the process is complicated because writes to the cache may still not have been flushed to memory at the point when a Write Back page is victimized. This is indicated by a `DIRTY` flag in the page's CPLB.

By default, Write Back pages are "clean", in that they do not have the `DIRTY` flag set. When a write occurs to a clean Write Back page, a Protection Violation exception is raised to indicate that the page is being written to. The `_cplb_mgr` routine flags the page's CPLB as `DIRTY`, and allows the write to continue. This time, it succeeds.

If, at some later time, the now-`DIRTY` CPLB is victimized, there may still be writes in the cache which have not yet been written to memory. The `_cplb_mgr` routine flushes any such writes from the cache back to memory before evicting the CPLB.

Because Write Through pages always update the memory location with the new cached value, Write Through pages do not need flushing when they are evicted. They also do not need to be marked as `DIRTY`. Consequently, Write Through pages do not trigger an exception on first write to the page.

Because of the additional overheads of flushing Write Back pages, the victimization process chooses victim CPLBs in the following order of preference:

1. Unused (for example, invalid) CPLBs
2. Unlocked CPLBs that are Write Through or Write Back and Clean
3. Unlocked CPLBs that are Dirty Write Back

Note that only unlocked CPLBs are selected as victims. Locked CPLBs are never selected. In particular, it is necessary to ensure that the CPLB management routines reside in pages which are covered by locked CPLBs, to prevent the CPLB management routines from evicting themselves or their configuration data.

To assist in this, the CPLB management routines reside in the `cplb_code` section; their configuration data resides in the `cplb_data` section. Both of these sections may be explicitly mapped to locked areas within the `.LDF` file.

Cache Flushing

When the `_cplb_mgr` routine has to flush modified Write Back data to memory, it uses the `flush_data_cache` routine. This routine can also be called explicitly, if circumstances require it. The routine searches all active pages for valid, modified pages that are cached in Write Back mode, and flushes their contents back to memory. This is a time-consuming process, and is dependent on the size of the modified data page:

Page Size	Approximate cost of flushing
1K	400 instructions
4K	1500 instructions
1M	6000 instructions
4M	6000 instructions

These costs are only approximate, because they only take into account the number of instructions executed, and do not include the costs of data transfers from cache to external memory. The actual cost is greatly influenced by the amount of modified data residing in the caches.

If it necessary to ensure smaller areas of memory are flushed to memory, the `flush_data_buffer` routine may be used:

```
#include <cplbtab.h>
void flush_data_buffer(void *start, void *end, int invalidate);
```

This routine flushes back to memory any changes in the data cache that apply to the address range specified by the `start` and `end` parameters. If the `invalidate` parameter is non-zero, the routine also invalidates the data cache for the address range, so that the next access to the range will require a re-fetch from memory.

Using the _cplb_mgr Routine

The `_cplb_mgr` routine is intended to be invoked by the application's exception handler. A minimal exception handler is installed by the default startup code, and its source can be found within your VisualDSP installation in the file `Blackfin\lib\src\libc\crt\cplbhdr.s`.

Typically, the exception handler delegates CPLB Misses and Protection Violations to `_cplb_mgr`, and handles all other exceptions itself. The `_cplb_mgr` routine is defined as (in C nomenclature):

```
int cplb_mgr(int code, int cplb_ctrl);
```

Code indicates the kind of exception raised, such as:

Code Value	Meaning
0	Instruction CPLB Miss
1	Data CPLB Miss
2	Protection Violation (assumed to be First-write to Write Back Data page)

The routine accepts the current value of ___cplb_ctrl as the second parameter.

There are several error codes that _cplb_mgr can return, defined in <cplb.h> as:

Return Code	Meaning
CPLB_RELOADED	Successfully updated CPLB table
CPLB_NO_UNLOCKED	All CPLBs are locked, so cannot be evicted. This indicates that the CPLBs in the configuration table are badly configured, as this should never occur.
CPLB_NO_ADDR_MATCH	The address being accessed, that triggered the exception, is not covered by any of the CPLBs in the configuration table. The application is presumably misbehaving.
CPLB_PROT_VIOL	The address being accessed, that triggered the exception, is not a first-write to a clean Write Back Data page, and so presumably is a genuine violation of the page's protection attributes. The application is misbehaving.

If _cplb_mgr returns an error indicator, the exception handler must decide how to handle the error. The default exception handler installed by the startup code delegates each of these failure conditions—plus one other—to handler functions in the run-time library.

These functions are:

```
void _unknown_exception_occurred(void);
void _cplb_miss_all_locked(void);
void _cplb_miss_without_replacement(void);
void _cplb_protection_violation(void);
```

These functions are stubs that can be replaced for comprehensive error-handling. They enter an infinite loop with verbose labels indicating the kind of error that has occurred.

Caching and Asynchronous Change

Care must be taken when using the cache in systems with asynchronous change. There are two levels of asynchronous data change:

1. Data that may change beyond the scope of the current thread, but within the scope of the system. This includes variables which may be updated by other threads in the system (if using a multi-threaded architecture). This kind of data must be marked volatile, so that the compiler knows not to store local copies in registers (but may be located in cached memory), since all threads access the data through the cache.

2. Data that may change beyond the scope of the cache as well as beyond the scope of the current thread. This includes memory-mapped registers (which cannot be cached) and data in memory which is updated by external means, such as DMA transfers, or host/target file I/O. Such data must be marked as volatile, so that the compiler knows not to keep copies in registers. This data may not be placed in cached memory since the cache does not see the change and provides date copies to the application. Alternatively, the cache copy must be invalidated before accessing memory, in case it has been updated.

Migrating LDFs From Previous VisualDSP++ Installations

The default `.LDF` files in VisualDSP++4.0 have undergone a number of changes in order to make caching simpler with minimal configuration for the user. The `.LDF` files derived from earlier versions of VisualDSP++ can benefit from these improvements by making a few modifications.

Link CPLB Configuration Tables Directly
Prior to VisualDSP++4.0, the default CPLB configuration tables were generic across platforms, and could be retrieved directly from the library. The default tables now cover each processor's memory map more compre-

hensively, and so each processor has its own `cplbtabn.doj` object file. This object file should be linked directly into the application, either by adding it to the project, or by including it in the `$OBJECTS` definition line.

For example,

```
$OBJECTS = CRT, $COMMAND_LINE_OBJECTS ,cplbtab533.doj ENDCRT;
```

Add Guard Symbols

The guard symbols defined by the `.LDF` files are optional, but greatly ease detection of mismatched `LDF/___cplb_ctrl` configurations. If the LDF does not define these guard symbols, default versions from the library are used, which indicate that all cache areas are available for use (which was the default assumption prior to VisualDSP++4.0).

C/C++ Preprocessor Features

Several features of the C/C++ preprocessor are used by VisualDSP++ to control the programming environment. The `ccblkfn` compiler provides standard preprocessor functionality, as described in any C text. The following extensions to standard C are also supported:

`//` end of line (C++ style) commands

`#warning` directive

For more information about these extensions, see "Preprocessor-Generated Warnings" on page 1-195 and "C++ Style Comments" on page 1-102. For ways to write macros, refer to "Writing Preprocessor Macros" on page 1-228.

This section contains:

- "Predefined Macros"
- "Writing Preprocessor Macros" on page 1-228

Predefined Macros

The `ccblkfn` compiler defines a number of macros to provide information about the compiler, source file, and options specified. These macros can be tested, using the `#ifdef` and related directives, to support your program's needs. Similar tailoring is done in the system header files.

For the list of predefined assertions, see "-A name(tokens)" on page 1-25.

Macros such as `__DATE__` can be useful if incorporated into the text strings. The # operator within a macro body is useful in converting such symbols into text constructs.

Table 1-15 describes the predefined compiler macros.

Table 1-15. Predefined Compiler Macro Listing

Macro	Function
__ADSPBLACKFIN__	Always defines __ADSPBLACKFIN__ as 1
__ADSPLPBLACKFIN__	Defines __ADSPLPBLACKFIN__ as 1 when the target processor (set using the -proc switch) is one of low-power core parts. These include ADSP-BF531, ADSP-BF532, ADSP-BF533, ADSP-BF534, ADSP-BF536, ADSP-BF537, ADSP-BF538, ADSP-BF539, ADSP-BF561 or ADSP-BF566 processors
__ANALOG_EXTENSIONS__	Defines __ANALOG_EXTENSIONS__ as 1, unless you compile with the -pedantic or -pedantic-errors switches
__cplusplus	Defines __cplusplus to be 199711L when you compile in C++ mode
__DATE__	The preprocessor expands this macro into the preprocessing date as a string constant. The date string constant takes the form mm dd yyyy (ANSI standard).
__DOUBLES_ARE_FLOATS__	The macro __DOUBLES_ARE_FLOATS__ is defined to 1 when the size of the double type is the same as the single precision float type. When the compiler switch -double-size-64 is used, the macro is not defined.
__ECC__	Always defines __ECC__ as 1
__EDG__	Always defines __EDG__ as 1. This definition signifies that an Edison Design Group front end is being used.
__EDG_VERSION__	Always defines __EDG_VERSION__ as an integral value representing the version of the compiler's front end
__EXCEPTIONS	Defines __EXCEPTIONS to be 1 when C++ exception handling is enabled (using the -eh switch)
__FILE__	The preprocessor expands this macro into the current input file name as a string constant. The string matches the name of the file specified on the command line or in a preprocessor #include command (ANSI standard).
__LANGUAGE_C	Always defines __LANGUAGE_C as 1; present when C compiler calls use it to specify headers

Table 1-15. Predefined Compiler Macro Listing (Cont'd)

Macro	Function
__LINE__	The preprocessor expands this macro into the current input line number as a decimal integer constant (ANSI standard).
__NO_BUILTIN	Defines __NO_BUILTIN as 1 when you compile with the -no-builtin command-line switch
__RTTI	Defines __RTTI to be 1 when C++ run-time type information is enabled (using the -rtti switch).
__SIGNED_CHARS__	Defines __SIGNED_CHARS__ as 1, unless you compile with the -unsigned-char command-line switch.
__STDC__	Always defines __STDC__ as 1
__STDC_VERSION__	Always defines __STDC_VERSION__ as 199409L
__TIME__	The preprocessor expands this macro into the preprocessing time as a string constant. The date string constant takes the form hh:mm:ss (ANSI standard).
__VERSION__	Defines __VERSION__ as a string constant giving the version number of the compiler used to compile this module.
__VERSIONNUM__	Defines __VERSIONNUM__ as a numeric variant of __VERSION__ constructed from the version number of the compiler. Eight bits are used for each component in the version number and the most significant byte of the value represents the most significant version component. As an example, a compiler with version 7.1.0.0 defines __VERSIONNUM__ as 0x07000100 and 7.1.1.10 would define __VERSIONNUM__ to be 0x0701010A.

Writing Preprocessor Macros

A macro is a name standing for a block of text that the preprocessor substitutes for. Use the `#define` preprocessor command to create a macro definition. When a macro definition has arguments, the block of text the preprocessor substitutes can vary with each new set of arguments.

Multi-statement Macros

Whenever possible, use inline functions rather than multi-statement macros. If multi-statement macros are necessary, define such macros to allow invocation like function calls. This will make your source code easier to read and maintain.

The following two code segments define two versions of the macro `SKIP_SPACES`.

```
/* SKIP_SPACES, regular macro */
#define SKIP_SPACES (p, limit)   \{
   char *lim = (limit); \
      while ((p) != lim)           { \
         if (*(p)++ != ' ')      { \
            (p)--; \
            break; \
         } \
      } \
}
/* SKIP_SPACES, enclosed macro */
#define SKIP_SPACES (p, limit)  \
   do { \
      char *lim = (limit);          \
      while ((p) != lim)         { \
         if (*(p)++ != ' ')      { \
            (p)--;   \
            break;   \
         }   \
      }   \
} while (0)
```

Enclosing the first definition within the `do {…} while (0)` pair changes the macro from expanding to a compound statement to expanding to a single statement. With the macro expanding to a compound statement, you would sometimes need to omit the semicolon after the macro call in order to have a legal program. This leads to a need to remember whether a function or macro is being invoked for each call and whether the macro needs a trailing semicolon or not.

With the `do {…} while (0)` construct, you can treat the macro as a function and put the semicolon after it.

For example,

```
/* SKIP_SPACES, enclosed macro, ends without ';' */
if (*p != 0)
   SKIP_SPACES (p, lim);
else …
```

This expands to

```
if (*p != 0)
   do {
      …
   } while (0);
else …
```

Without the `do {…} while (0)` construct, the expansion would be:

```
if (*p != 0)
   {
      ...
   };   /* Probably not intended syntax */
else
```

C/C++ Run-Time Model and Environment

This section provides a full description of the Blackfin processor run-time model and run-time environment. The run-time model, which applies to compiler-generated code, includes descriptions of layout of the stack, data access, and call/entry sequence. The C/C++ run-time environment includes the conventions that C/C++ routines must follow to run on Blackfin processors. Assembly routines linked to C/C++ routines must follow these conventions.

Analog Devices recommends that assembly programmers maintain stack conventions.

Figure 1-5 provides an overview of the run-time environment issues that must be considered when writing assembly routines that link with C/C++ routines including the "C/C++ Run-Time Header and Startup Code" on page 1-232. The run-time environment issues include the following items.

- Memory usage conventions

  "Using Memory Sections" on page 1-243

  "Using Multiple Heaps" on page 1-245

  "Using Data Storage Formats" on page 1-261

- Register usage conventions

  "Dedicated Registers" on page 1-251

  "Call Preserved Registers" on page 1-252

  "Scratch Registers" on page 1-252

  "Stack Registers" on page 1-253

- Program control conventions

"Managing the Stack" on page 1-253

"Transferring Function Arguments and Return Value" on page 1-257

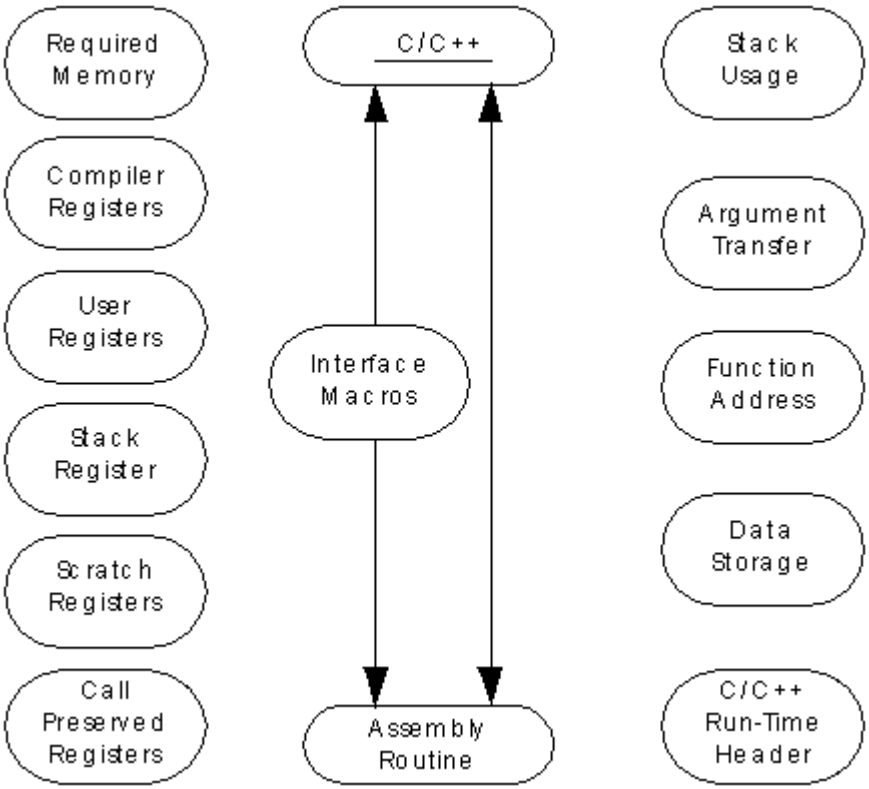

Figure 1-5. Assembly Language Interfacing Overview

C/C++ Run-Time Header and Startup Code

The C/C++ run-time (CRT) header is code that is executed after the processor jumps to the start address on reset. The CRT header sets the machine into a known state and calls _main. CRT code can be included in a project in one of the following ways:

- The **Project Wizard** can be used to automatically generate a customised CRT in a project. Refer to the *VisualDSP++ 4.0 User's Guide* for further information.
- The macro `USER_CRT` can be defined at link-time to specify a custom user-defined CRT object is to be included in the project build.
- Default CRT objects are provided for all platforms in the run-time libraries, and are linked against for all C/C++ projects if the link-time macro `USER_CRT` is not defined.

This section contains:

- "CRT Header Overview"
- "CRT Description" on page 1-234

CRT Header Overview

The CRT header (called "CRT" in this section):

- Sets up required system services, such as stack and event-handling
- Disables system services that power up in an undefined state, such as circular buffers and hardware loops
- Initializes standard library facilities, such as standard input/output (`stdio`)
- Invokes C++ constructors, as required

The CRT ensures that when execution enters `_main`, the processor's state obeys the C Application Binary Interface (ABI), and that global data declared by the application have been initialized as required by the C/C++ standards. It arranges things so that `_main` appears to be "just another function" invoked by the normal function invocation procedure.

Not all applications require the same configuration. For example, C++ constructors are invoked only for applications that contain C++ code. The list of optional configuration items is long enough that determining whether to invoke each one in turn at runtime would be overly costly.

For this reason, the **Project Wizard** allows a CRT to be generated which includes the minimum amount of code necessary given the user-selected options. Additionally, the pre-built CRTs are supplied in several different configurations, which can be specified at link-time via LDF macros.

The CRT header is used for projects that use C, C++, and VDK. Assembly language projects do not provide a default run-time header; you must provide your own.

The source assembly file for the pre-compiled CRTs is located under the VisualDSP++ installation directory, in the file `basiccrt.s`, in the directory `Blackfin\lib\src\libc\crt`. Each of the pre-built CRT objects are built from this default CRT source. The different configurations are produced by the definition of various macros (refer to EE-238 for more information).

The list of operations performed by the CRT (startup code) can include (not necessarily in the following order):

- Setting registers to known/required values
- Disabling hardware loops
- Disabling circular buffers
- Setting up default event handlers and enabling interrupts

- Initializing the Stack and Frame Pointers
- Enabling the cycle counter
- Configuring the memory ports used by the two DAGs
- Setting the processor clock speed
- Copying data from the flash memory to RAM
- Initializing device drivers
- Setting up memory protection and caches
- Changing processor interrupt priority
- Initializing profiling support
- Invoking C++ constructors
- Invoking `_main`, with supplied parameters
- Invoking `_exit` on termination

What the CRT Does Not Do

The CRT does not initialize actual memory hardware. However, some properties of external SDRAM may be configured if the CRT contains code to customize clock and power settings. The initialization of the external SDRAM is left to the boot loader because it is possible (and even likely) that the CRT itself will need to be moved into external memory before being executed.

CRT Description

The following sections describe the main operations which may be performed by the CRT, dependent on the selected **Project Wizard** options, or which of the pre-built CRTs is included in the build.

Declarations

The CRT begins with a number of preprocessor directives that "include" the appropriate platform-definition header and set up a few constants:

- `IVBl` and `IVBh` give the address of the Event Vector Table.
- `UNASSIGNED_VAL` is a bit pattern that indicates that the register/memory location has not yet been written to by the application. See "Mark Registers" on page 1-241 and "Terminate Stack Frame Chain" on page 1-242.
- `INTERRUPT_BITS` is the default interrupt mask. By default, it enables the lowest-priority interrupt, `IVG15`. This default mask can also be overridden at runtime by your own version of `__install_default_handlers`; see "Event Vector Table" on page 1-236 for details.
- For some platforms, `SYSCFG_VALUE` is the initialization value for the System Configuration register (`SYSCFG`).
- Within the standard CRT source, `SET_CLOCK_SPEED` determines whether the CRT will attempt to improve the processor's default clock speed. Some accompanying definitions follow. More advanced processor clock and power setting is available via the **Project Wizard** options. See "Clock Speed" on page 1-238.

Start and Register Settings

The CRT declares its first code label as `start`. This required label is referenced by `.LDF` files, which explicitly resolve this label to the processor's reset address.

First, the CRT disables facilities that could be enabled on start-up, due to their random power-up states, as follows:

- SYSCFG is set to SYSCFG_VALUE, according to anomaly 05-00-0109 for ADSP-BF531, ADSP-BF532, ADSP-BF533, and ADSP-BF561 processors.
- Hardware loops are disabled. This prevents the jump-back-to-loop-start behavior should the "loop bottom" register correspond to the start of an instruction.
- Circular buffer lengths are set to zero. The CRT makes use of the Iregs and calls functions that may use them. Furthermore, the C/C++ ABI requires that circular buffers are disabled on entry to (and exit from) compiled functions, so the circular buffers must be disabled before invoking _main.

Event Vector Table

The reset vector (fixed) and emulation events (not touched by the C ABI), are not defined by the CRT. The processor's lowest-priority event, IVG15, is set to point to supervisor_mode, a label that appears later in the CRT and is used to facilitate the switch to supervisor mode. The remaining entries of the Event Vector Table are loaded with the address of the __unknown_exception_occurred dummy event handler, which results in defined behavior to aid debugging.

Additionally, if caching or memory protection is enabled (either selected via the **Project Wizard**, or configured by the user-defined value of the ___cplb_ctrl variable), an exception handler is required to process possible events raised by the memory system. Therefore, the default handler, ___cplb_hdr, is installed into the exception entry of the Event Vector Table.

For details on ___cplb_ctrl, refer to "Caching and Memory Protection" on page 1-208.

You may install additional handlers; for your convenience, the CRT calls a function to do this. The function, `__install_default_handlers`, is an empty stub, which you may replace with your own function that installs additional or alternative handlers, before the CRT enables events. The function's C prototype is:

```
short _install_default_handlers(short mask);
```

The CRT passes the default enable mask, (`INTERRUPT_BITS`), as a parameter, and considers the return value to be an updated enable mask. If you install additional handlers, you must return an updated enable mask to reflect this.

See the *VisualDSP++ 4.0 VDK User's Guide* for details on how to configure ISRs for applications that use VDK.

Stack Pointer and Frame Pointer

The Stack Pointer (SP) is set to point to the top of the stack, as defined in the `.LDF` file by the symbol `ldf_stack_end`. Specifically, the Stack Pointer is set to point just past the top of the stack. Because stack pushes are pre-decrement operations, the first push moves the Stack Pointer so that it refers to the actual stack top.

The User Stack Pointer (USP) and Frame Pointer (FP) are set to point to the same address.

Twelve bytes are then claimed from the stack. This is because the C ABI requires callers to allocate stack space for the parameters of callees, and that all functions require at least twelve bytes of stack space for registers `R0-R2`. Therefore, the CRT claims these twelve bytes as the incoming parameters for functions called before invoking `_main`.

Cycle Counter

The CRT enables the cycle counter, so that the `CYCLES` and `CYCLES2` registers are updated. This is not necessary for general program operation, but is desirable for measuring performance.

DAG Port Selection

For ADSP-BF531, ADSP-BF532, ADSP-BF533, ADSP-BF534, ADSP-BF536, ADSP-BF537, ADSP-BF538, ADSP-BF539, ADSP-BF561 and ADSP-BF566 processors, the CRT configures the DAGs to use different ports for accessing memory. This reduces stalls when the DAGs issue memory accesses in parallel.

Clock Speed

The **Project Wizard** allows clock and power configuration to be specified for supported hardware. This generated CRT then contains calls to the System Services library Power Management module to configure the processor clock and power settings. For further information, see the *Device Drivers and System Services Library Manual*.

Alternatively, within the default CRT source, if `SET_CLOCK_SPEED` is set is set to one, the CRT changes the default clock speed from approximately 300 MHz to approximately 600 MHz, provided the silicon is of a suitable revision. This is determined by reading the `DSPID` memory-mapped register and ensuring that the silicon version is 0.2 or greater.

To change the clock speed, the CRT uses two library routines with the following C prototypes:

```
int pll_set_system_vco(int msel, int df, int lockcnt)
int pll_set_system_clocks(int csel, int ssel)
```

These two routines return zero for success and negative for error, but the CRT does not test their return values.

The EZ-KIT Lite boards use a 27 MHz on-board clock. The CRT uses the following constants as parameters to these routines to set the PLL and clocks:

VCO = 27*0x16 = 594 MHz

CCLK = VCO = 594 MHz

SCLK = VCO/5 = 118 MHz

Memory Initialization

Memory initialization is a two-stage process:

1. At linktime, the Memory Initializer utility processes the `.DXE` file to generate a table of code and data memory areas that must be initialized during booting.
2. At runtime, when the application starts, the run-time library function `_mi_initialize` processes the table to copy the code and data from the Flash device to volatile memory.

If the application has not been processed by the Memory Initializer, or if the Memory Initializer did not find any code or data that required such movement, the `_mi_initialize` function returns immediately. The generated CRT includes a call to `_mi_initialize` if the "*Run-time memory initialization*" option is selected in the **Project Wizard**. The default CRT source always includes the call.

The CRT does not enable external memory. The configuration of physical memory hardware is the responsibility of the boot loader and must be complete before the CRT is invoked.

Device Initialization

The CRT calls `_init_devtab` to initialize device drivers that support `stdio`. This routine:

1. Initializes the internal file tables
2. Invokes the `init` routine for each device driver registered at build-time
3. Associates `stdin`, `stdout`, and `stderr` with the default device driver

Device initialization is included in the generated CRT if the **C/C++ I/O and I/O device support** option of the **Project Wizard** is selected (default). For information on the device drivers supported by `stdio`, refer to "Extending I/O Support To New Devices" on page 3-35.

CPLB Initialization

If instruction or data Cache Protection Lookaside Buffers (CPLBs) are to be enabled, the CRT calls `_cplb_init` to do this, passing the value of `___cplb_ctrl` as a parameter. For details on CPLB initialization, refer to "Caching and Memory Protection" on page 1-208.

Enable Interrupts

If instruction or data Cache Protection Lookaside Buffers (CPLBs) are to be enabled, the CRT calls function `_cplb_init` to do this, passing the value of `___cplb_ctrl` as a parameter. The declaration and initialization of the global `___cplb_ctrl` variable is included in the generated CRT if memory protection or caching is selected via the **Project Wizard**. Otherwise, the default library definition of the variable is used, unless overridden by a declaration in user code. For details on CPLB initialization, refer to "Caching and Memory Protection" on page 1-208.

Lower Processor Priority

The CRT lowers the process priority to the lowest supervisor mode level (IVG15). It first raises IVG15 as an event, but this event cannot be serviced while the processor remains at the higher priority level of Reset.

The CRT sets RETI to be the still_interrupt_in_ipend label, at which there is an RTI instruction, and is the next instruction to be executed. This results in all bits representing interrupts of a higher priority than IVG15 being cleared. In normal circumstances this would include only the reset interrupt, but occasionally this may not be the case, for example, if the program is restarted during an ISR.

The pending IVG15 interrupt is now allowed to proceed, and the handler set up earlier in the CRT (at the label supervisor_mode) is executed. Thus, execution flows from the "return" from Reset level to the supervisor_mode label, while changing processor mode from the highest supervisor level to the lowest supervisor level.

If other events are enabled (memory system exceptions or other events installed via your own version of the default handlers stub), they could be taken between the return from Reset and entering IVG15. Therefore, the remaining parts of the CRT may not execute when event handlers are triggered.

The CRT's first action after entering IVG15 is to re-enable the interrupt system so that other higher-priority interrupts can be processed.

Mark Registers

The UNASSIGNED_FILL value is written into R2-R7 registers and P0-P5 registers if the **Project Wizard** option's "**Initialize data registers to a known value**" is selected (or the **UNASSIGNED_FILL** macro is defined when re-building the default CRT source).

Terminate Stack Frame Chain

Each stack frame is pointed to by the Frame Pointer and contains the previous values of the Frame Pointer and RETS. The CRT pushes two instances of `UNASSIGNED_VAL` onto the stack, indicating that there are no further active frames. The C++ exception support library uses these markers to determine whether it has walked back through all active functions without finding one with a catch for the thrown exception.

Again, the CRT allocates twelve bytes for outgoing parameters of functions that will be called from the CRT.

Profiler Initialization

If the **Project Wizard** option's "**Enable Profiling**" is selected, the CRT initializes the instrumented-code profiling library by calling `monstartup`. This routine zeroes all counters and ensures that no profiling frames are active. The instrumented-code profiling library uses `stdio` routines to write the accumulated profile data to `stdout` or to a file.

Instrumented-code profiling is specified with the `-p`, `-p1`, and `-p2` compiler switches (see "-p[1|2]" on page 1-47). These are added to the compilation options if necessary by the **Project** Wizard. If any of the object files was compiled to include this profiling, the prelinker detects this and sets link-time macros which select a profiler-enabled pre-built CRT object if the **Project Wizard** is not in use.

C++ Constructor Invocation

The `___ctorloop` function runs all of the global-scope constructors, and is always called from the **Project Wizard**-generated CRT, and from the C++ enabled pre-built CRTs (which the `.LDF` file selects if a C++ compiled object has been detected).

Each global-scope constructor has a pointer to it located in the `ctor` section. When all of these sections are gathered together during linking, they form a table of function pointers. The `___ctorloop` library function walks

the length of this table, invoking each constructor in turn. The .LDF file arranges for the special object crtn.doj to be linked after the CRT; it contains a terminator for the ctor table.

Argument Parsing

The __getargv function is called to parse any provided arguments (normally an empty list) into the __Argv global array. This function returns the number of arguments found which, along with __Argv, form the argc and argv parameters for _main. Within the default CRT source, if FIOCRT is not defined, argc is set to zero and argv is set to an empty list, statically defined within the CRT.

Calling _main and _exit

The _main is called, using the argc and argv just defined. Embedded programs are not expected to return from _main, but many legacy and non-embedded programs do. Therefore, the return value from _main is immediately passed to _exit to gracefully terminate the application. _exit is not expected to return.

Using Memory Sections

The C/C++ run-time environment requires that a specific set of memory section names are used for placing code in memory. In assembly language files, these names are used as labels for the .SECTION directive. In the .LDF file, these names are used as labels for the output section names within the SECTIONS{} command. For information on the LDF syntax and other information on the linker, see the *VisualDSP++ 4.0 Linker and Utilities Manual*.

Code Storage

The code section, program, is where the compiler puts all the program instructions that it generates when compiling the program. The cplb_code

section exists so that memory protection management routines can be placed into sections of memory that are always configured as being available.

Data Storage
The data section, `data1`, is where the compiler puts global and static data in memory. The data section, `constdata`, is where the compiler puts data that has been declared as "`const`". By default, the compiler places global zero-initialized data into a "BSS-style" section, called "`bsz`", unless the compiler is invoked with the `-no-bss` option (on page 1-41). The section `voldata` exists to contain volatile data that must be separated from cached memory, but the compiler does not automatically place volatile data into the `voldata` section. The `cplb_data` section exists so that configuration tables used to manage memory protection can be placed in memory areas that are always flagged as accessible.

Run-Time Stack
The run-time stack sections, `stack` and `sysstack`, are where the compiler puts the run-time stack in memory. The processor starts off in supervisor mode, which uses `sysstack` for its run-time stack. You may choose to link your application so the processor is switched to user mode during startup (see "Startup Code Overview" on page 1-196). When in user mode, the run-time stack is in the `stack` section. When linking, use your `.LDF` file to map this section. Because the run-time environment cannot function without this section, you must define it.

The run-time stack is a 32-bit wide structure, growing from high memory to low memory. The compiler uses the run-time stack as the storage area for local variables and return addresses. See "Managing the Stack" on page 1-253 for more information.

Run-Time Heap Storage
The run-time heap section, `heap`, is where the compiler puts the run-time heap in memory. When linking, use your `.LDF` file to map the heap section. To dynamically allocate and deallocate memory at runtime, the C run-time library includes four functions:

```
malloc()   calloc()   realloc()   free()
```

These functions allocate from the `heap` section of memory by default. The `.LDF` file must define `ldf_heap_space`, `ldf_heap_end` and `ldf_heap_length` to allow the `_Sbrk()` function to determine the location of the heap. Default values for these are defined in the default `.LDF` file.

Using Multiple Heaps

The Blackfin C/C++ run-time library supports the standard heap management functions `calloc`, `free`, `malloc`, and `realloc`. By default, there is a single heap, called the default heap, which serves all allocation requests that do not explicitly specify an alternative heap. The default heap is defined in the standard Linker Description File and the run-time header.

The code written by the user can define any number of additional heaps. These heaps serve allocation requests that are explicitly directed to them. These additional heaps can be accessed via the extension routines `heap_calloc`, `heap_free`, `heap_malloc` and `heap_realloc`.

Multiple heaps allow the programmer to serve allocations using fast-but-scarce memory or slower-but-plentiful memory as appropriate. This section describes how to define a heap, work with heaps, use the heap interface, and free space in the heap.

Defining a Heap

Heaps can be defined at link time or at runtime. In both cases, a heap has three attributes:

- Start (base) address (the lowest usable address in the heap)
- Length (in bytes)
- User identifier (`userid`, a number >= 1)

The default system heap, defined at link time, always has `userid 0`. In addition, heaps have indices This is like the `userid`, except that the index is assigned by the system. All the allocation and deallocation routines use heap indices, not heap user IDs; a `userid` can be converted to its index using `_heap_lookup()` (see "Defining Heaps at LinkTime"). Be sure to pass the correct identifier to each function.

Defining Heaps at LinkTime

Link-time heaps are defined in the `heaptab.s` file in the library, and their start address, length, and `userid` are held in three 32-bit words. The heaps are in a table called "`_heap_table`". This table must contain the default heap (`userid 0`) first and must be terminated by an entry that has a base address of zero.

The addresses placed into this table can be literal addresses, or they can be symbols that are resolved by the linker. The default heap uses symbols generated by the linker through the `.LDF` file.

The `_heap_table` table can live in constant memory. It is used to initialize the run-time heap structure, `___heaps`, when the first request to a heap is made. When allocating from any heap, the library initializes `___heaps` using the data in `_heap_table`, and sets `___nheaps` to be the number of available heaps.

There is a compiled-in upper limit on the number of heaps allowed. This is defined by `MAXHEAPS` in `heapinfo.h` and is currently set to 4. This is used purely to determine the size of the `___heaps` structure.

Because the allocation routines use heap indices instead of heap user IDs, a heap installed in this fashion needs to have its `userid` mapped into an index before it can be used explicitly:

```
int _heap_lookup(int userid);    // returns index
```

Defining Heaps at RunTime

Heaps may also be defined and installed at runtime, using the `_heap_install()` function:

```
int _heap_install(void *base, size_t length, int userid);
```

This function can take any section of memory and start using it as a heap. It returns the heap index allocated for the newly installed heap, or a negative value if there was some problem (see "Tips for Working With Heaps").

Reasons why `_heap_install()` may return an error status include, but are not limited to:

- Not enough space available in the `___heaps` table
- A heap using the specified `userid` already exists
- New heap appears too small to be usable (length too small)

Tips for Working With Heaps

Heaps may not start at address zero (`0x0000 0000`). This address is reserved and means "no memory could be allocated". It is the null pointer on the Blackfin platform.

Not all memory in a heap is available for use by the user. Some of the memory (a handful of words) is reserved for housekeeping. Thus, a heap of 256 bytes is unable to serve four blocks of 64 bytes.

Memory reserved for housekeeping precedes the allocated blocks. Thus, if a heap begins at `0x0800 0000`, this particular address is never returned to the user program as the result of an allocation request; the first request returns an address some way into the heap.

The base address of a heap must be appropriately aligned for an 8-byte memory access. This means that allocations can then be used for vector operations.

The lengths of heaps should be multiples of powers of two for most efficient space usage. The heap allocator works in block sizes such as 256, 512 or 1024 bytes.

Standard Heap Interface

The standard functions, `calloc` and `malloc`, always allocate a new object from the default heap. If `realloc` is called with a null pointer, it too allocates a new object from the default heap.

Previously allocated objects can be deallocated with the `free` or `realloc` functions. When a previously allocated object is resized with `realloc`, the returned object is always in the same heap as the original object.

The `space_unused` function returns the number of bytes unallocated in the heap with index 0. Note that you may not be able to allocate all of this space due to heap fragmentation and the overhead that each allocated block needs.

The `space_unused` function returns the number of bytes unallocated in the heap with index 0. Note that you may not be able to allocate all of this space due to heap fragmentation, and the overhead that each allocated block needs.

Using the Alternate Heap Interface

The C run-time library provides the alternate heap interface functions `heap_calloc`, `heap_free`, `heap_malloc`, and `heap_realloc`. These routines work in exactly the same way as the corresponding standard functions without the `heap_` prefix, except that they take an additional argument that specifies the heap index.

For example,

```
void *_heap_calloc(int idx, size_t nelem, size_t elsize)
void *_heap_free(int idx, void *)
void *_heap_malloc(int idx, size_t length)
void *_heap_realloc(int idx, void *, size_t length)
```

The actual entry point names for the alternate heap interface routines have an initial underscore. The `stdlib.h` standard header file defines macro equivalents without the leading underscores.

Note that for

```
heap_realloc(idx, NULL, length)
```

the operation is equivalent to

```
heap_malloc(idx, length)
```

However, for

```
heap_realloc(idx, ptr, length)
```

where `ptr != NULL`, the supplied `idx` parameter is ignored; the reallocation is always done from the heap that `ptr` was allocated from, even if a `memcpy` function is required within the heap.

Similarly,

```
heap_free(idx, ptr)
```

ignores the supplied index parameter, which is specified only for consistency—the space indicated by `ptr` is always returned to the heap from which it was allocated.

The `heap_space_unused(int idx)` function returns the number of bytes unallocated in the heap with index `idx`. The function returns `-1` if there is no heap with the requested heap index.

C++ Run-time Support for the Alternate Heap Interface

The C++ run-time library provides support for allocation and release of memory from an alternative heap via the new and delete operators.

Heaps should be initialized with the C run-time functions as described. These heaps can then be used via the new and delete mechanism by simply passing the heap ID to the new operator. There is no need to pass the heap ID to the delete operator as the information is not required when the memory is released.

The routines are used as in the example below:

```
#include <heapnew>

char *alloc_string(int size, int heapID)
{
     char *retVal = new(heapID) char[size];
     return retVal;
}

void free_string(char *aString)
{
     delete aString;
}
```

Freeing Space

When space is “freed”, it is not returned to the “system”. Instead, freed blocks are maintained on a free list within the heap in question. The blocks are coalesced where possible.

It is possible to re-initialize a heap, emptying the free list and returning all the space to the heap itself, such as

```
int _heap_init(int index)
```

This returns zero for success, and nonzero for failure. Note, however, that this discards all records within the heap, so they may not be used if there are any live allocations on the heap still outstanding.

Dedicated Registers

The C/C++ run-time environment specifies a set of registers whose contents should not be changed except in specific defined circumstances. If these registers are changed, their values must be saved and restored. The dedicated register values must always be valid for every function call (especially for library calls) and for any possible interrupt.

Dedicated registers are:

`SP (P6) — FP (P7)`

`L0 — L3`

- The `SP (P6)` and `FP (P7)` are the Stack Pointer and the Frame Pointer registers, respectively. The compiler requires that both are 4-byte aligned registers pointing to valid areas within the stack section.
- The `L0—L3` registers define the lengths of the DAG's circular buffers. The compiler makes use of the DAG registers, both in linear mode and in circular buffering mode. The compiler assumes that the Length registers are zero, both on entry to functions and on return from functions, and ensures this is the case when it generates calls or returns. Your application may modify the Length registers and make use of circular buffers, but you must ensure that the Length registers are appropriately reset when calling compiled functions, or returning to compiled functions. Interrupt handlers must store and restore the Length registers, if making use of DAG registers.

Call Preserved Registers

The C/C++ run-time environment specifies a set of registers whose contents must be saved and restored. Your assembly function must save these registers during the function's prologue and restore the registers as part of the function's epilogue. The call preserved registers must be saved and restored if they are modified within the assembly function; if a function does not change a particular register, it does not need to save and restore the register. The registers are:

P3—P5
R4—R7

Scratch Registers

The C/C++ run-time environment specifies a set of registers whose contents do not need to be saved and restored. Note that the contents of these registers are not preserved across function calls.

Scratch registers are:

P0	Used as the Aggregate Return Pointer
P1-P2	
R0-R3	The first three words of the argument list are always passed in R0, R1 and R2 if present (R3 is not used for parameters).
LB0-LB1	
LC0-LC1	
LT0-LT1	
ASTAT	Including CC
A0-A1	
I0-I3	
B0-B3	
M0-M3	

Loop Counters, Overlays and DMA'd Code
The compiler does not ensure that the loop counter registers `LC0` and `LC1` are zero on entry or exit from a function. This does not normally cause a problem because the exit point of a hardware loop is unique within the program, and the compiler ensures that the only path to the exit is through the corresponding loop setup instruction.

If overlays are being used, or if code is being DMA'd into faster memory for execution, this may no longer be the case. It is possible for an overlay or a DMA'd function to set up a loop that terminates at address A, and then for a different overlay or DMA'd function to have different code occupying address A at a later point in time. If a hardware loop is still active—`LC0` or `LC1` is non-zero—at the point when the instruction at address A is reached, then undefined behavior results as the hardware loop "jumps" back to the start of the loop.

Therefore, in such cases, it is necessary for the overlay manager or the DMA manager to reset loop counters to ensure no hardware loops remain active that might relate to the address range covered by the variant code.

Stack Registers

The C/C++ run-time environment reserves a set of registers for controlling the run-time stack. These registers may be modified for stack management, but must be saved and restored. Stack registers are:

`SP (P6)` – Stack pointer
`FP (P7)` – Frame pointer

Managing the Stack

The C/C++ run-time environment uses the run-time stack to store automatic variables and return addresses. The stack is managed by a Frame Pointer (FP) and a Stack Pointer (SP) and grows downward in memory, moving from higher to lower addresses.

A stack frame is a section of the stack used to hold information about the current context of the C/C++ program. Information in the frame includes local variables, compiler temporaries, and parameters for the next function.

The Frame Pointer serves as a base for accessing memory in the stack frame. Routines refer to locals, temporaries, and parameters by their offset from the Frame Pointer.

Figure 1-6 shows an example section of a run-time stack. In the figure, the currently executing routine, `Current()`, was called by `Previous()`, and `Current()` in turn calls `Next()`. The state of the stack is as if `Current()` has pushed all the arguments for `Next()` onto the stack and is just about to call `Next()`.

Stack usage for passing any or all of a function's arguments depends upon the number and types of parameters to the function.

As you write assembly routines, note that operations to restore stack and frame pointers are the responsibility of the called function.

To enter and perform a function, follow this sequence of steps:

- **Linking Stack Frames** – The return address and the caller's FP are saved on the stack, and FP set pointing to the beginning of the new (callee) stack frame. SP is decremented to allocate space for local variables and compiler temporaries.
- **Register Saving** – Any registers that the function needs to preserve are saved on the stack frame, and SP is set pointing to the top of the stack frame.

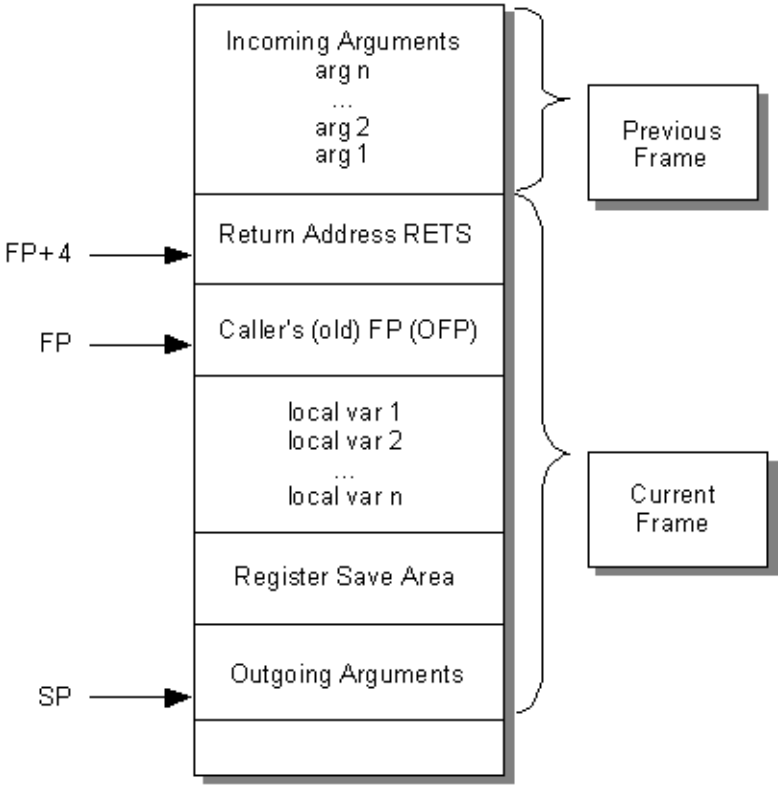

Figure 1-6. Example Run-Time Stack

At the end of the function, these steps must be performed:

- **Restore Registers** – Any registers that had been preserved are restored from the stack frame, and SP is set pointing to the top of the stack frame.
- **Unlinking Stack Frame** – The frame pointer is restored from the stack frame to the caller's FP, `RETS` is restored from the stack frame to the return address, and SP is set pointing to the top of the caller's stack frame.

A typical function prologue would be

```
LINK   16;
[--SP]=(R7:4);
SP + = —16;
[FP+8]=R0; [FP+12]=R1; [FP+16]=R2;
```

where

`LINK 16;`
is a special linkage instruction that saves the return address and the frame pointer, and updates the Stack Pointer to allocate 16 bytes for local variables.

`[--SP]=(R7:4);`
allocates space on the stack and saves the registers in the save area.

`SP + = —16;`
allocates space on the stack for outgoing arguments. Always allocate at least 12 bytes on the stack for outgoing arguments, even if the function being called requires less than this.

`[FP+8]=R0; [FP+12]=R1; [FP+16]=R2;`
saves the argument registers in the argument area.

A matching function epilogue would be

```
SP + = 16;
P0=[FP+4];
(R7:4)=[SP++];
UNLINK;
JUMP (P0);
```

where

`SP + = 16;`
reclaims the space on the stack that was used for outgoing arguments.

`P0=[FP+4]`
loads the return address into register `P0`.

`(R7:4)=[SP++];`
restores the registers from the save area and reclaims the area.

`UNLINK;`
is a special instruction that restores the frame pointer and stack pointer.

`JUMP (P0);`
returns to the caller.

The section "Transferring Function Arguments and Return Value" provides additional details on function call requirements.

Transferring Function Arguments and Return Value

The C/C++ run-time environment uses a set of registers and the run-time stack to transfer function parameters to assembly routines. Your assembly language functions must follow these conventions when they call (or when called by) C/C++ functions.

Passing Arguments

The details of argument passing are most easily understood in terms of a conceptual argument list. This is a list of words on the stack. Double arguments are placed starting on the next available word in the list, as are structures. Each argument appears in the argument list exactly as it would in storage, and each separate argument begins on a word boundary.

The actual argument list is like the conceptual argument list except that the contents of the first three words are placed in registers `R0`, `R1` and `R2`. Normally this means that the first three arguments (if they are integers or pointers) are passed in registers `R0` to `R2` with any additional arguments being passed on the stack.

If any argument is greater than one word, it occupies multiple registers. The caller is responsible for extending any `char` or `short` arguments to 32-bit values.

When calling a C function, at least twelve bytes of stack space must be allocated for the function's arguments, corresponding to `R0–R2`. This applies even for functions that have less that 12 bytes of argument data, or that have fewer than three arguments.

The details of argument passing do not change for variable argument lists. For example, if a function is declared as:

```
int varying(char *fmt, ...) { /* ... */ }
```

it may receive one or more arguments. As with other functions, the first argument, `fmt`, is passed in `R0`, and other arguments are passed in `R1`, and then `R2`, and then on the stack, as required.

Variable argument lists are processed using the macros defined in the `stdarg.h` header file. The `va_start()` function obtains a pointer to the list of arguments which may be passed to other functions, or which may be walked by the `va_arg()` macro.

To support this, the compiler begins variable argument functions by flushing `R0`, `R1` and `R2` to their reserved spaces on the stack:

```
_varying:
     [SP+0] = R0;
     [SP+4] = R1;
     [SP+8] = R2;
```

The `va_start()` function can then take the address of the last non-varying argument (`fmt`, in the example above, at `[SP+0]`), and `va_arg()` can walk through the complete argument list on the stack.

Return Values

If a function returns a `short` or a `char`, the "callee" is responsible for effecting any sign or zero extension that is needed.

- For functions returning aggregate values occupying less than or equal to 32 bits, the result is returned in `R0`.
- For aggregate values occupying greater than 32 bits, and less than or equal to 64 bits, the result is returned in register pair `R0`, `R1`.
- For functions returning aggregate values occupying more than 64 bits, the caller allocates the return value object on the stack and the address of this object is passed to the callee as a hidden argument in register `P0`.

The callee must copy the return value into the object at the address in `P0`. Table 1-16 provides examples of passed parameters.

Table 1-16. Examples of Parameter Passing

Function Prototype	Parameters Passed as	Return Location
`int test(int a, int b, int c)`	`a in R0, b in R1, c in R2`	`in R0`
`char test(int a, char b, char c)`	`a in R0, b in R1, c in R2`	`in R0`
`int test(int a)`	`a in R0`	`in R0`
`int test(char a, char b, char c, char d, char e)`	`a in R0, b in R1, c in R2, d in [FP+20], e in [FP+24]`	`in R0`
`int test(struct *a, int b, int c)`	`a (addr) in R0, b in R1, c in R2`	`in R0`

Table 1-16. Examples of Parameter Passing (Cont'd)

struct s2a { char ta; char ub; int vc;} int test(struct s2a x, int b, int c)	x.ta and x.ub in R0, x.vc in R1, b in R2, c in [FP+20]	in R0
struct foo *test(int a, int b, int c)	a in R0, b in R1, c in R2	(address) in R0
void qsort(void *base, int nel, int width, int (*compare)(const void *, const void *))	base(addr) in R0, nel in R1, width in R2, compare(addr) in [FP+20]	

struct s2 {
char t;
char u;
int v;
}
struct s2 test(int a,
int b, int c)
	a in R0,
	b in R1,
	c in R2
		in R0 (s.t and s.u)
		and in R1 (s.v)

struct s3 {
char t;
char u;
int v;
int w;
}
struct s3 test(int a,
int b, int c)
	a in R0,
	b in R1,
	c in R2
		in *P0 (based on value
		of P0 at the call, not
		necessarily at the
		return)

Using Data Storage Formats

The sizes of intrinsic C/C++ data types are selected by Analog Devices so that normal C/C++ programs execute with hardware-native data types, and, therefore, at high speed. The C/C++ run-time environment uses the intrinsic C/C++ data types and data formats that appear in Table 1-17 and are shown in Figure 1-7 and Figure 1-8.

Table 1-17. Data Storage Formats and Data Type Sizes

Type	Bit Size	Number Representation	`sizeof` returns
`char`	8 bits signed	8-bit two's complement	1
`unsigned char`	8 bits unsigned	8-bit unsigned magnitude	1
`short`	16 bits signed	16-bit two's complement	2
`unsigned short`	16 bits unsigned	16-bit unsigned magnitude	2
`int`	32 bits signed	32-bit two's complement	4
`unsigned int`	32 bits unsigned	32-bit unsigned magnitude	4
`long`	32 bits signed	32-bit two's complement	4
`unsigned long`	32 bits unsigned	32-bit unsigned magnitude	4
`long long`	64 bits signed	64-bit two's complement	8
`unsigned long long`	64 bits unsigned	64-bit unsigned magnitude	8
`pointer`	32 bits	32-bit two's complement	4
`function pointer`	32 bits	32-bit two's complement	4
`double`	32 bits	32-bit IEEE single-precision	4
`float`	32 bits	32-bit IEEE single-precision	4
`double`	64 bits	64-bit IEEE double-precision	8
`long double`	64 bits	64-bit IEEE	8
`fract16`	16 bits signed	1.15 fract	2
`fract32`	32 bits signed	1.31 fract	4

The floating-point and 64-bit data types are implemented using software emulation, so must be expected to run more slowly than hard-supported native data types. The emulated data types are `float`, `double`, `long double`, `long long` and `unsigned long long`.

`fract16` and `fract32` are not actually intrinsic data types—they are `typedefs` to `short` and `long`, respectively. In C, you need to use built-in functions to do basic arithmetic (see "Fractional Value Built-in Functions in C" on page 1-104). You cannot do `fract16*fract16` and get the right result. In C++, for `fract` data, the classes "`fract`" and "`shortfract`" define the basic arithmetic operators.

Floating-Point Data Size

On Blackfin processors, the `float` data type is 32 bits, and the `double` data type default size is 32 bits. This size is chosen because it is the most efficient. The 64-bit `long double` data type is available if more precision is needed, although this is more costly because the type exceeds the data sizes supported natively by hardware.

In the C language, floating-point literal constants default to `double` data type. When operations involve both `float` and `double`, the `float` operands are promoted to `double` and the operation is done at `double` size. By having `double` default to a 32-bit data type, the Blackfin compiler usually avoids additional expense during these promotions. This does not however fully conform to the C and C++ standards which require that the `double` type supports at least 10 digits of precision.

The `-double-size-64` switch (on page 1-29) sets the size of the `double` type to 64 bits if the additional precision, or full standard conformance, is required.

The `-double-size-64` switch causes the compiler to treat the `double` data type as a 64-bit data type, instead of a 32-bit data type. This means that all values are promoted to 64 bits, and consequently incur more storage and cycles during computation. The switch does not affect the size of the `float` data type, which remains at 32 bits.

Consider the following case.

```
float add_two(float x) { return x + 2.0; }   // has promotion
```

When compiling this function, the compiler promotes the `float` value `x` to `double`, to match the literal constant `2.0`. The addition becomes a `double` operation, and the result is truncated back to a `float` before being returned.

By default, or with the `-double-size-32` switch (on page 1-29), the promotion and truncation operations are empty operations—they require no work because the `float` and `double` types default to the same size. Thus, there is no cost.

With the `-double-size-64` switch, the promotion and truncation operations do require work because the `double` constant 2.0 is a 64-bit value. The `x` value is promoted to 64 bits, a 64-bit addition is performed, and the result is truncated to 32 bits before being returned.

In contrast,

```
float add_two(float x) { return x + 2.0f; }   // no promotion
```

Because the literal constant `2.0f` has a "`f`" suffix, it is a float-type constant, not a double-type constant. Thus, both operands to the addition are `float` type, and no promotion or truncation is necessary. This version of the function does not produce any performance degradation when the `-double-size-64` switch is used.

You must be consistent in your use of the `-double-size-{32|64}` switch.

Consider the two files, such as:

```
file x.c:
double add_nums(double x, double y) { return x + y; }

file y.c:
extern double add_nums(double, double);
double times_two(double val) { return add_nums(val, val); }
```

Both files must be compiled with the same usage of `-double-size{32|64}`. Otherwise, `times_two()` and `add_nums()` will be exchanging data in mismatched formats, and incorrect behavior will occur. Table 1-18 shows the results for the various permutations:

Table 1-18. Use of the -double-size-{32|64} Switch

x.c	y.c	Result
default	default	Okay
default	-double-size-32	Okay
-double-size-32	default	Okay
-double-size-32	-double-size-32	Okay
-double-size-64	-double-size-64	Okay
-double-size-32	-double-size-64	Error
-double-size-64	-double-size-32	Error

If a file does not make use of any double-typed data, it may be compiled with the `-double-size-any` switch (on page 1-29), to indicate this fact. Files compiled in this way may be linked with files compiled with `-double-size-32` or with `-double-size-64`, without conflict.

Conflicts are detected by the linker and result in linker error 1151, “Input sections have inconsistent qualifiers”.

Floating-Point Binary Formats

The Blackfin compiler supports IEEE and non-IEEE floating-point formats.

IEEE Floating-Point Format

By default, the Blackfin compiler provides floating point emulation using IEEE single- and double-precision formats. Single-precision IEEE format (Figure 1-7) provides a 32-bit value, with 23 bits for mantissa, 8 bits for exponent, and 1 bit for sign. This format is used for the `float` data type, and for the double data type by default and when the `-double-size-32` switch is used.

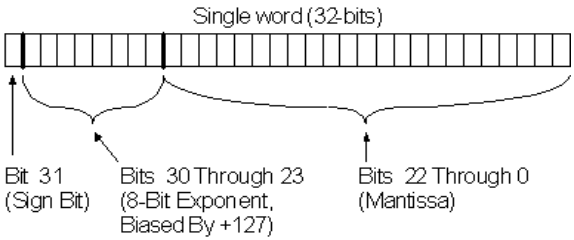

The single word (32-Bit) data storage format equates to:

$$-1^{Sign} \times 1.Mantissa \times 2^{(Exponent - 127)}$$

Where:

Sign	Comes from the sign bit
Mantissa	Represents the fractional part of the Mantissa (23-Bits). The 1. is assumed in this format
Exponent	Represents the 8-Bit exponent

Figure 1-7. Data Storage Format for Float and Double Types

Double-precision IEEE format (Figure 1-8) provides a 64-bit value, with 52 bits for mantissa, 11 bits for exponent, and 1 bit for sign. This format is used for the `long double` data type, and for the `double` data type when the `-double-size-64` switch is used.

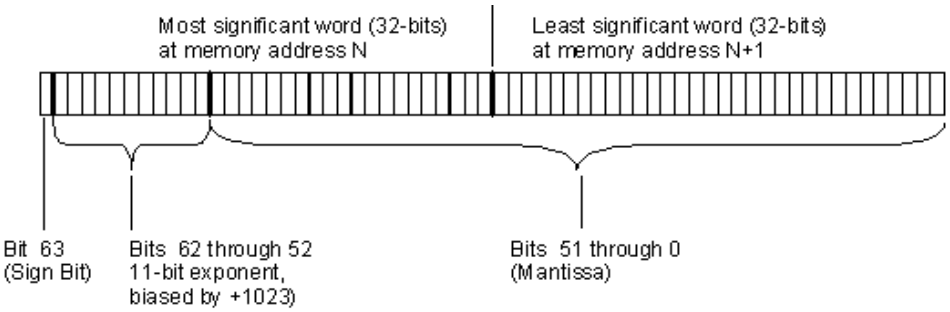

The two word (64-bit) data storage format equates to:

$$-1^{Sign} \times 1.Mantissa \times 2^{(Exponent - 1023)}$$

Where:

Sign	Comes from the Sign Bit
Mantissa	Represents the fractional part of the Mantissa (52-bits). The 1. is assumed in this format.
Exponent	Represents 11-bit exponent

Figure 1-8. Double-Precision IEEE Format

Variants of IEEE Floating-Point Support

The Blackfin compiler supports two variants of IEEE floating-point support. These variants are implemented in terms of two alternative emulation libraries, selected at link time.

The two alternative emulation libraries are:

- The default IEEE floating-point library
  It is a high-performance variant, which relaxes some of the IEEE rules in the interests of performance. This library assumes that its inputs will be value numbers, rather than `Not-a-number` values. This library can also explicitly be selecting via the `-fast-fp` switch (on page 1-31).
- An alternative IEEE floating-point library
  It is a strictly-conforming variant, which offers less performance, but includes all the input-checking that has been relaxed in the default library. The strictly-conforming library can be selected via the `-ieee-fp` switch (on page 1-36).

The default `.LDF` file links in the appropriate archive(s) depending on the setting of the link-time macro `IEEEFP`. If the `-ieee-fp` switch has been specified, the compiler defines the macro and the LDF links the application against the non-default, IEEE-conforming library. Conversely, if the link-time macro `IEEEFP` is not defined, then the default `.LDF` file arranges for the application to be linked against the default, high-performance, floating-point archives.

Fract16 and Fract32 Data Representation

The `fract16` type represents a single 16-bit signed fractional value, while the `fract32` type represents a 32-bit signed fractional value. Both types have the same range, [-1.0,+1.0). However, `fract32` has twice the precision. They are not intrinsic data storage formats, they are simply `typedefs`.

```
typedef short fract16;
typedef long  fract32;
```

The `fract` data representation is shown in Figure 1-9

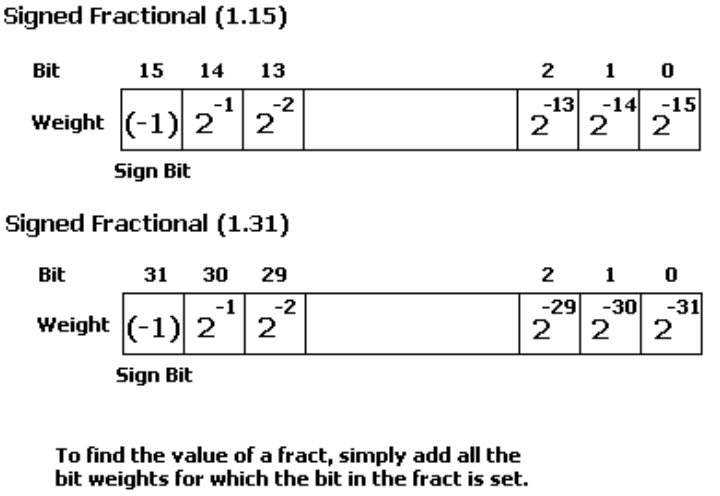

Figure 1-9. Data Storage Format for fract16 and fract32

Therefore, to represent 0.25 in `fract16`, the HEX representation would be `0x2000` (2^{-2}). For -0.25 in `fract32`, the HEX would be `0xe000 0000` ($-1+2^{-1}+2^{-2}$)). For -1, the HEX representation in `fract16` would be `0x8000` (-1). `fract16` and `fract32` cannot represent +1 exactly, but they get quite close with `0x7fff` for `fract16`, or `0x7fff ffff` for `fract32`.

There is also a `fract2x16` data type, which is two `fract16s` packed into 32 bits. The first two bytes belong to one `fract16`, and the second two bytes belong to the other. There are also built-in functions that work with `fract2x16` parameters.

C/C++ and Assembly Interface

This section describes how to call assembly language subroutines from within C/C++ programs, and how to call C/C++ functions from within assembly language programs. Before attempting to perform either of these operations, familiarize yourself with the information about the C/C++ run-time model (including details about the stack, data types, and how arguments are handled) in "C/C++ Run-Time Model and Environment" on page 1-230. At the end of this reference, a series of examples demonstrate how to mix C/C++ and assembly code.

This section describes:

- "Calling Assembly Subroutines From C/C++ Programs"
- "Calling C/C++ Functions From Assembly Programs" on page 1-272

Calling Assembly Subroutines From C/C++ Programs

Before calling an assembly language subroutine from a C/C++ program, create a prototype to define the arguments for the assembly language subroutine and the interface from the C/C++ program to the assembly language subroutine. Even though it is legal to use a function without a prototype in C/C++, prototypes are a strongly-recommended practice for good software engineering. When the prototype is omitted, the compiler cannot perform argument-type checking and assumes that the return value is of type integer and uses K&R promotion rules instead of ANSI promotion rules.

The compiler prefaces the name of any external entry point with an underscore. Therefore, declare your assembly language subroutine's name with a leading underscore.

The run-time model defines some registers as *scratch* registers and others as *preserved* or *dedicated* registers. Scratch registers can be used within the assembly language program without worrying about their previous contents. If more room is needed (or an existing code is used) and you wish to use the preserved registers, you *must save* their contents and then *restore* those contents before returning.

Use the dedicated or stack registers for their intended purpose only; the compiler, libraries, debugger, and interrupt routines depend on having a stack available as defined by those registers.

The compiler also assumes the machine state does not change during execution of the assembly language subroutine.

Do not change any machine modes (for example, certain registers may be used to indicate circular buffering when those register values are nonzero).

If arguments are on the stack, they are addressed via an offset from the stack pointer or frame pointer. A good way to explore how arguments are passed between a C/C++ program and an assembly language subroutine is to write a dummy function in C/C++ and compile it using the `save temporary files` option (the `-save-temps` command-line switch).

The following example includes the global volatile variable assignments to indicate where the arguments can be found upon entry to `asmfunc`.

```
// Sample file for exploring compiler interface …
// global variables … assign arguments there just so
// we can track which registers were used
// (type of each variable corresponds to one of arguments):

int global_a;
float global_b;
int * global_p;

// the function itself:

int asmfunc(int a, float b, int * p)
```

```
{
   // do some assignments so assembly file file will show
       where args are:
   global_a = a;
   global_b = b;
   global_p = p;

   // value gets loaded into the return register:
   return 12345;
}
```

When compiled with the `-save-temps` switch being set, the following code is produced.

```
   .section program;
   .align 2;
_asmfunc:
       P0.L = .epcbss;
       P0.H = .epcbss;
       [P0+ 0] = R0;
       R0 = 0x1234 (X);
       [P0+ 4] = R1;
       [P0+ 8] = R2;
       RTS;

._asmfunc.end:
   .global _asmfunc;
   .type _asmfunc,STT_FUNC;

   .section data1;

    .align 8;
.epcbss:
      .byte _global_a[4];
      .global _global_a;
      .type _global_a,STT_OBJECT;
       .byte _global_b[4];
       .global _global_b;
       .type _global_b,STT_OBJECT;
       .byte _global_p[4];
       .global _global_p;
       .type _global_p,STT_OBJECT;
.epcbss.end:
```

Calling C/C++ Functions From Assembly Programs

You may want to call a C/C++ callable library and other functions from within an assembly language program. As discussed in "Calling Assembly Subroutines From C/C++ Programs" on page 1-269, you may want to create a test function to do this in C/C++, and then use the code generated by the compiler as a reference when creating your assembly language program and the argument setup. Using volatile global variables may help clarify the essential code in your test function.

The run-time model defines some registers as *scratch* registers and others as *preserved* or *dedicated*. The contents of the scratch registers may be changed without warning by the called C/C++ function. If the assembly language program needs the contents of any of those registers, you *must save* their contents before the call to the C/C++ function and then *restore* those contents after returning from the call.

Use the dedicated registers for their intended purpose only; the compiler, libraries, debugger, and interrupt routines all depend on having a stack available as defined by those registers.

Preserved registers can be used; their contents are not changed by calling a C/C++ function. The function always saves and restores the contents of preserved registers if they are going to change.

If arguments are on the stack, they are addressed via an offset from the stack pointer or frame pointer. Explore how arguments are passed between an assembly language program and a function by writing a dummy function in C/C++ and compiling it with the `save temporary files` option (see the `-save-temps` switch on page 1-54). By examining the contents of volatile global variables in a `*.s` file, you can determine how the C/C++ function passes arguments, and then duplicate that argument setup process in the assembly language program.

The stack must be set up correctly before calling a C/C++ callable function. If you call other functions, maintaining the basic stack model also facilitates the use of the debugger. The easiest way to do this is to define a C/C++ main program to initialize the run-time system; maintain the stack until it is needed by the C/C++ function being called from the assembly language program; and then continue to maintain that stack until it is needed to call back into C/C++. However, make sure the dedicated registers are correct. You do not need to set the FP prior to the call; the caller's FP is never used by the recipient.

Using Mixed C/C++ and Assembly Naming Conventions

A user should be able to use C/C++ symbols (function or variable names) in assembly routines and use assembly symbols in C/C++ code. This section describes how to name and use C/C++ and assembly symbols.

To name an assembly symbol that corresponds to a C/C++ symbol, add an underscore prefix to the C/C++ symbol name when declaring the symbol in assembly. For example, the C/C++ symbol `main` becomes the assembly symbol `_main`.

To use a C/C++ function or variable in an assembly routine, declare it as global in the C program. Import the symbol into the assembly routine by declaring the symbol with the `.EXTERN` assembler directive.

To use an assembly function or variable in your C/C++ program, declare the symbol with the `.GLOBAL` assembler directive in the assembly routine and import the symbol by declaring the symbol as `extern` in the C program.

Table 1-19 shows several examples of the C/C++ and assembly interface naming conventions.

Table 1-19. C/C++ Naming Conventions for Symbols

In the C/C++ Program	In the Assembly Subroutine
`int c_var; /*declared global*/`	`.extern _c_var;` `.type _c_var,STT_OBJECT;`
`void c_func(void);`	`.global _c_func;` `.type _c_func,STT_FUNC;`

`extern int asm_var;`	`.global _asm_var;`
	`.type _asm_var,STT_OBJECT;`
	`.byte = 0x00,0x00,0x00,0x00`

`extern void asm_func(void);`	`.global _asm_func;`
	`.type _asm_func,STT_FUNC;`
	`_asm_func:`

Compiler C++ Template Support

The compiler provides template support C++ templates as defined in the ISO/IEC 14882:1998 C++ standard, with the exception that the export keyword is not supported.

Template Instantiation

Templates are instantiated automatically during compilation using a linker feedback mechanism. This involves compiling files, determining any required template instantiations, and then recompiling those files making the appropriate instantiations. The process repeats until all required instantiations have been made. Multiple recompilations may be required in the case when a template instantiation is made that requires another template instantiation to be made.

By default, the compiler uses a method called *implicit instantiation*, which is common practice, and results in having both the specification and definition available at point of instantiation. This involves placing template specifications in a header (for example, "`.h`") file and the definitions in a source (for example, "`.cpp`") file. Any file being compiled that includes a header file containing template specifications will instruct the compiler to implicitly include the corresponding "`.cpp`" file containing the definitions of the compiler.

For example, you may have the header file "`tp.h`"

```
template <typename A> void func(A var)
```

and source file "`tp.cpp`"

```
template <typename A> void func(A var)
{
...code...
}
```

Two files "file1.cpp" and "file2.cpp" that include "tp.h" will have file "tp.cpp" included implicitly to make the template definitions available to the compilation.

If there is a desire not to use the implicit inclusion method then the switch -no-implicit-inclusion should be passed to the compiler. In the example, we have been discussing, "tp.cpp" will then be treated as a normal source file and should be explicitly linked into the final product.

Regardless of whether implicit instantiation is used or not, the compilation process involves compiling one or more source files and generating a ".ti" file corresponding to the source files being compiled. These ".ti" files are then used by the prelinker to determine the templates to be instantiated. The prelinker creates a ".ii" file and recompiles one or more of the files instantiating the required templates.

The prelinker ensures that only one instantiation of a particular template is generated across all objects. For example, the prelinker ensures that if both "file1.cpp" and "file2.cpp" invoked the template function with an int, that the resulting instantiation would be generated in just one of the objects.

Identifying Uninstantiated Templates

If for some reason the prelinker is unable to instantiate all the templates that are required for a particular link then a link error will occur. For example:

```
[Error li1021] The following symbols referenced in processor 'P0'
could not be resolved:
    'Complex<T1> Complex<T1>::_conjugate() const [with T1=short]
[_conjugate__16Complex__tm__2_sCFv_18Complex__tm__4_Z1Z]' refer-
enced from '.\Debug\main.doj'
    'T1 *Buffer<T1>::_getAddress() const [with T1=Complex<short>]
[_getAddress__33Buffer__tm__19_16Complex__tm__2_sCFv_PZ1Z]'
referenced from '.\Debug\main.doj'
    'T1 Complex<T1>::_getReal() const [with T1=short]
```

```
[_getReal__16Complex__tm__2_sCFv_Z1Z]' referenced from
'.\Debug\main.doj'

Linker finished with 1 error
```

Careful examination of the linker errors reveals which instantiations have not been made. Below are some examples.

```
Missing instantiation:
   Complex<short> Complex<short>::conjugate()
Linker Text:
   'Complex<T1> Complex<T1>::_conjugate() const [with T1=short]
   [_conjugate__16Complex__tm__2_sCFv_18Complex__tm__4_Z1Z]'
   referenced from '.\Debug\main.doj'

Missing instantiation:
   Complex<short> *Buffer<Complex<short>>::getAddress()
Linker Text:
   'T1 *Buffer<T1>::_getAddress() const [with T1=Complex<short>]
   [_getAddress__33Buffer__tm__19_16Complex__tm__2_sCFv_PZ1Z]'
   referenced from '.\Debug\main.doj'

Missing instantiation:
   Short Complex<short>::getReal()
Linker Text:
   'T1 Complex<T1>::_getReal() const [with T1=short]
   [_getReal__16Complex__tm__2_sCFv_Z1Z]' referenced from
   '.\Debug\main.doj'
```

There could be many reasons for the prelinker being unable to instantiate these templates, but the most common is that the `.ti` and `.ii` files associated with an object file have been removed. Only source files that can contain instantiated templates will have associated `.ti` and `.ii` files, and without this information, the prelinker may not be able to complete its task. Removing the object file and recompiling will normally fix this problem.

Another possible reason for uninstantiated templates at link time is when implicit inclusion (described above) is disabled but the source code has been written to require it. Explicitly compiling the `.cpp` files that would normally have been implicitly included and adding them to the final link is normally all that is needed to fix this.

Another likely reason for seeing the linker errors above is invoking the linker directly. It is the compiler's responsibility to instantiate C++ templates, and this is done automatically if the final link is performed via the compiler driver. The linker itself contains no support for instantiating templates.

2 ACHIEVING OPTIMAL PERFORMANCE FROM C/C++ SOURCE CODE

This chapter provides guidance on tuning your application to achieve the best possible code from the compiler. Some implementation choices are available when coding an algorithm, and understanding their impact is crucial to attaining optimal performance.

This chapter contains:

- "General Guidelines" on page 2-3
- "Improving Conditional Code" on page 2-24
- "Loop Guidelines" on page 2-25
- "Using Built-In Functions in Code Optimization" on page 2-35
- "Smaller Applications: Optimizing for Code Size" on page 2-40
- "Using Pragmas for Optimization" on page 2-42
- "Useful Optimization Switches" on page 2-51
- "Example: How the Optimizer Works" on page 2-52
- "Assembly Optimizer Annotations" on page 2-54

The focus of this chapter is obtaining maximal code performance from the compiler. Most of these guidelines also apply when optimizing for minimum code size, although some techniques specific to that goal are also discussed.

The first section looks at some general principles, and how the compiler can lend the most help to your optimization effort. Optimal coding styles are then considered in detail. Special features such as compiler switches, built-in functions, and pragmas are also discussed. The chapter ends with a short example to demonstrate how the optimizer works.

Small examples are included throughout this chapter to demonstrate key points. Some show recommended coding styles, others identify styles to be avoided or code that it may be possible to improve. These are marked as "**Good**" and "**Bad**", respectively.

General Guidelines

It is useful to bear in mind the following basic strategy when writing an application:

1. Try to choose an algorithm that is suited to the architecture being targeted. For example, there may be a trade-off between memory usage and algorithm complexity that may be influenced by the target architecture.
2. Code the algorithm in a simple, high-level generic form. Keep the target in mind, especially with respect to choices of data types.
3. You can then turn your attention towards code tuning. For critical code sections, think more carefully about the strengths of the target platform, and make any non-portable changes where necessary.

Choose the language as appropriate.

As the programmer, your first decision is to choose whether to implement your application in C or C++. This decision may be influenced by performance considerations. C++ code using only C features has very similar performance to a pure C source. Many higher level C++ features (for example, those resolved at compilation, such as namespaces, overloaded functions and inheritance) have no performance cost. However, use of some other features may degrade performance, and so the performance loss must be weighed against the richness of expression available in C++. Examples of features that may degrade performance include virtual functions or classes used to implement basic data types.

This section contains:

- "How the Compiler Can Help" on page 2-4
- "Data Types" on page 2-8
- "Getting the Most From IPA" on page 2-13

- "Indexed Arrays Versus Pointers" on page 2-18
- "Function Inlining" on page 2-19
- "Using Inline asm Statements" on page 2-20
- "Memory Usage" on page 2-21

How the Compiler Can Help

The compiler provides many facilities designed to help the programmer including compiler optimizer, statistical profiler, PGO, and IPA optimizers.

Using the Compiler Optimizer

There is a vast difference in performance between code compiled optimized and code compiled non-optimized. In some cases, optimized code can run ten or twenty times faster. Optimization should always be used when measuring performance or shipping code as product.

The optimizer in the C/C++ compiler is designed to generate efficient code from source that has been written in a straightforward manner. The basic strategy for tuning a program is to present the algorithm in a way that gives the optimizer excellent visibility of the operations and data, and hence the greatest freedom to safely manipulate the code. Future releases of the compiler will continue to enhance the optimizer, and expressing algorithms simply will provide the best chance of benefiting from such enhancements.

Note that the default setting (or "debug" mode within the VisualDSP++ IDDE) is for non-optimized compilation in order to assist programmers in diagnosing problems with their initial coding. The optimizer is enabled in VisualDSP++ by checking the **Enable optimization** check box under the **Project Options ->Compile** tab. This adds the `-O` (enable optimization) switch (on page 1-45) to the compiler invocation. A "release" build from within VisualDSP++ automatically enables optimization.

Using the Statistical Profiler

Tuning an application begins with an understanding of which areas of the application are most frequently executed and therefore where improvements would provide the largest gains. The statistical profiling feature provided in VisualDSP++ is an excellent means for finding these areas. More details about how to use it may be found in the *VisualDSP++ 4.0 User's Guide*.

The particular advantage of statistical profiling is that it is completely unobtrusive. Other forms of profiling insert instrumentation into the code, disturbing the original optimization, code size, and register allocation to some degree.

The best methodology is usually to compile with both optimization and debug information generation enabled. In this way, you can obtain a profile of the optimized code while retaining function names and line number information. This gives you accurate results that correspond directly to the C/C++ source. Note that the compiler optimizer may have moved code between lines.

You can obtain a more accurate view of your application if you build it optimized but without debug information generation. You then obtain statistics that relate directly to the assembly code. The only problem with doing this may be in relating assembly lines to the original source. Do not strip out function names when linking, since keeping function names means you can scroll through the assembly window to instructions of interest.

In very complicated code, you can locate the exact source lines by counting the loops, unless they are unrolled. Looking at the line numbers in the assembly file (use the `-save-temps` switch to retain compiler generated assembly files, which will have the `.s` filename extension) may also help. The compiler optimizer may have moved code around so that it does not appear in the same order as in your original source.

Using Profile-Guided Optimization

Profile-Guided Optimization (PGO) is an excellent way to tune the compiler's optimization strategy for the typical run-time behavior of a program. There are many program characteristics that cannot be known statically at compilation time but can be provided by means of PGO. The compiler can use this knowledge to bring about benefits, such as more accurate branch prediction, improved loop transformations, and reduced code size. The technique is most relevant where the behavior of the application over different data sets is expected to be very similar.

The application is first compiled with the `-pguide` switch (on page 1-50) to create an executable file containing the necessary instrumentation for gathering profile data. Optimization should also be enabled—for best results, use the `-O` (on page 1-45) or `-ipa` (on page 1-36) switches.

The next step is to gather the profile; at present this may only be done using a simulator. To do this, run the executable with one or more training data sets. These training data sets should be representative of the data that you expect the application to process in the field. Note that unrepresentative training data sets can cause performance degradations when the application is used on real data. The profile is accumulated in a file with the extension `.pgo`. The final stage involves re-compiling the application using this gathered profile data. To do this, simply place the `.pgo` file on the command line. Optimization should also be enabled at this stage.

Note that when a `.pgo` file is placed on the command line, the `-O` optimization switch by default tries to balance between code performance and code-size considerations. More specifically, it is equivalent to using `-Ov 50`. To optimize solely for performance while using PGO, the switch `-Ov 100` should be used, or the optimization slider bar (located under the **Compiler** tab of the **Project Options** dialog box in VisualDSP++) moved up to its highest setting. The `-Ov n` switch is discussed further when looking at the specific issues involved in optimization for space (see "Smaller Applications: Optimizing for Code Size" on page 2-40).

PGO should always be performed as the very last optimization step. If the application source code is changed after gathering profile data, this profile data becomes invalid. The compiler does not use profile data when it can detect that it is inaccurate. However, it is possible to change source code in a way that is not detectable to the compiler (for example, by changing constants), and it is the programmer's job to ensure that the profile data being used for optimization remains accurate.

For more details on PGO, refer to "Optimization Control" on page 1-71.

Using Interprocedural Optimization

To obtain the best performance, the optimizer often requires information that can only be determined by looking outside the function that it is working on. For example, it helps to know what data can be referenced by pointer parameters, or whether a variable actually has a constant value. The `-ipa` compiler switch (on page 1-36) enables interprocedural analysis (IPA), which can make this data available. When this switch is used, the compiler is called again from the link phase to recompile the program using additional information obtained during previous compilations.

Because it only operates at link time, the effects of IPA are not seen if you compile with the `-S` switch (on page 1-53). To see the assembly file when IPA is enabled, use the `-save-temps` switch (on page 1-54), and look at the `.s` file produced after your program has been built.

As an alternative to IPA, you can achieve many of the same benefits by adding pragma directives and other declarations such as `__builtin_aligned` to provide information to the compiler about how each function interacts with the rest of the program.

These directives are further described in "Using __builtin_aligned" on page 2-15 and "Using Pragmas for Optimization" on page 2-42.

Data Types

The compiler directly supports the following scalar data types.

Single-Word Fixed-Point Data Types: Native Arithmetic	
`char`	8-bit signed integer
`unsigned char`	8-bit unsigned integer
`short`	16-bit signed integer
`unsigned short`	16-bit unsigned integer
`int`	32-bit signed integer
`unsigned int`	32-bit unsigned integer
`long`	32-bit signed integer
`unsigned long`	32-bit unsigned integer
Double-Word Fixed-Point Data Types: Emulated Arithmetic	
`long`	32-bit signed integer
`unsigned long`	32-bit unsigned integer
`long long`	64-bit signed integer
`unsigned long long`	64-bit unsigned integer
Floating-Point Data Types: Emulated Arithmetic	
`float`	32-bit float
`double`	The size of the `double` type differs depending on the compiler switches used. If the -double-size-64 switch is specified, the type represents a 64-bit `float`, otherwise it represents a 32-bit `float`.
`long double`	64-bit float

Fractional data types are represented using the integer types. Manipulation of these data types is best done by using the built-in functions, described in "System Support Built-In Functions" on page 2-37.

Optimizing a Struct

Memory can be saved, with no loss of performance, if a struct is declared with the smallest members first, progressing to the largest members. For example,

```
struct optimal_struct {
   char element1,element2;
   short element3;
   int element4;
};
```

will occupy 8 bytes of memory, whereas:

```
struct non_optimal_struct {
   char element1;      /* 3 bytes of padding */
   int element2;
   short element3;
   char element4;      /* 1 byte of padding  */
};
```

will occupy 12 bytes of memory.

When the compiler generates a memory access, the access will be to a 1-, 2- or 4-byte unit. Such accesses must be naturally aligned, meaning that 2-byte accesses must be to even addresses, and 4-byte accesses must be to addresses on a 4-byte boundary. Failure to align addresses results in a misaligned memory access exception.

The compiler is required to retain the order of members of a `struct`, and must ensure these alignment constraints are met. Therefore, by default, the compiler inserts any necessary padding to ensure that elements are aligned on their required boundaries. Padding is also inserted after the last member of a struct if required, to ensure that the struct's size is a multiple of the struct's strictest member alignment.

That is:

- If any member has a 4-byte alignment, the `struct` is a multiple of 4 bytes in size.
- Otherwise, if any member has a two-byte alignment, the `struct` is a multiple of two bytes in size.
- Otherwise, no end-of-struct padding is required.

Therefore, for a concrete example, if you have

```
struct non_optimal_struct test[2];
```

and if the compiler did not insert padding into the struct `non_optimal_struct`, the size of struct `non_optimal_struct` would be 8 bytes, and `test[]` array would be 16 bytes in size. Then, if

```
int x = test[1].element2;
```

this would be attempting to read an `int` (4 bytes) from a misaligned address (address of test+9), and thus a hardware exception (misaligned access) would occur.

Because the compiler adds appropriate padding in the struct `non_optimal_struct`, the `int` read will read a 4-byte aligned address (address of test+16), and the access will succeed.

As a rule of thumb, to get the smallest possible struct, elements should be placed in the struct in the following order:

```
typedef struct efficient_struct{
   size_1_elements a,...;
   size_2_elements b,...;
   size_4_or_greater_elements c,...;
}
```

The compiler supports greater density of structs through the use of the `#pragma pack(n)` directive. This allows you to reduce the necessary padding required in structs without reordering the struct's members. There is a trade-off implied, because the compiler must still observe the architecture's address-alignment constraints. When `#pragma pack(n)` is used, it means that a struct member is being accessed across the required alignment boundary, and the compiler must decompose the member into smaller, appropriately-aligned components and issue multiple accesses.

See "#pragma pack (alignopt)" on page 1-153 for more details.

Bit Fields

Use of bit fields in code can reduce the amount of data storage required by an application, but will normally increase the amount of code for an application (and thus make the application slower). This is because more code is needed to access a bit field than to access an intrinsic type (`char`, `int`, and so on). Also bit fields may prevent the compiler from performing optimizations that it could do on intrinsic types. However, depending on the use of bit fields, the total data bytes plus total code bytes may be less when using bit fields instead of intrinsic types. For example,

```
struct bitf {
   int item1:5;
   int item2:3;
   char item3;
   short item4;
};
```

This `struct` packs a a 5-bit item, a 3-bit item, an 8-bit item and a 16-bit item into 4 bytes.

The array `struct bitf arr[1000]` would save a significant amount of data space over a non-bit field version. However, compared to not using a bit field, more code would be generated to access the bit field members of the struct, and that code would be slower.

Avoiding Emulated Arithmetic

Arithmetic operations for some data types are implemented by library functions because the processor hardware does not directly support these types. Consequently, operations for these data types are far slower than native operations—sometimes by a factor of a hundred—and also produce larger code. These types are marked as "Emulated Arithmetic" in "Data Types".

The hardware does not provide direct support for division, so division and modulus operations are almost always multi-cycle operations, even on integral type inputs. If the compiler has to issue a full-division operation, it usually needs to generate a call to a library function. One notable situation in which a library call is avoided is for integer division when the divisor is a compile-time constant and is a power of two. In that case the compiler generates a shift instruction. Even then, a few fix-up instructions are needed after the shift, if the types are signed. If you have a signed division by a power of two, consider whether you can change it to unsigned in order to obtain a single-instruction operation.

When the compiler has to generate a call to a library function for one of these arithmetic operators that are not supported by the hardware, performance would suffer not only because the operation takes multiple cycles, but also because the effectiveness of the compiler optimizer is reduced.

For example, such an operation in a loop can prevent the compiler from making use of efficient zero-overhead hardware loop instructions. Also, calling the library to perform the required operation can change values held in scratch registers before the call. Therefore, the compiler has to generate more stores and loads from the data stack to keep values required after the call returns. Emulated arithmetic operators should therefore be avoided where possible, especially in loops.

Getting the Most From IPA

Interprocedural analysis (IPA) is designed to try to propagate information about the program to parts of the optimizer that can use it. This section looks at what information is useful, and how to structure your code to make this information easily accessible for analysis.

Initialize Constants Statically

IPA identifies variables that have only one value and replace them with constants, resulting in a host of benefits for the optimizer's analysis. For this to happen a variable must have a single value throughout the program. If the variable is statically initialized to zero, as all global variables are by default, and is subsequently assigned some other value at another point in the program, then the analysis sees two values and does not consider the variable to have a constant value.

For example,

```
#include <stdio.h>
int val;           // initialized to zero
void init() {
   val = 3;        // re-assigned
}
void func() {
   printf("val %d",val);
}
int main() {
   init();
   func();
}
```

Bad: IPA cannot see that `val` is a constant.

The code is better written as

```
   #include <stdio.h>
const int val = 3;    // initialized once
```

```
void init() {
}
void func() {
   printf("val %d",val);
}
int main() {
   init();
   func();
}
```

Good: IPA knows `val` is 3.

Word-Aligning Your Data

To make most efficient use of the hardware, it must be kept fed with data. In many algorithms, the balance of data accesses to computations is such that, to keep the hardware fully utilized, data must be fetched with 32-bit loads.

The hardware requires that references to memory be naturally aligned. Thus, 16-bit references must be at even address locations, and 32-bit references at word-aligned addresses. Therefore, for the most efficient code to be generated, ensure that data buffers are word-aligned.

The compiler helps to establish the alignment of array data. Top-level arrays are allocated at word-aligned addresses, regardless of their data types. To do this for local arrays means that the compiler also ensures that stack frames are kept word-aligned. However, arrays within structures are not aligned beyond the required alignment for their type. It may be worth using the `#pragma align n` directive to force the alignment of arrays in this case.

If you write programs (that only pass the address of the first element of an array as a parameter) and loops (that process these input arrays an element at a time, starting at element zero), then IPA should be able to establish that the alignment is suitable for full-width accesses.

Where an inner loop processes a single row of a multi-dimensional array, try to ensure that each row begins on a word boundary. In particular, two-dimensional arrays should be defined in a single block of memory rather than as an array of pointers to rows all separately allocated with `malloc`. It is difficult for the compiler to keep track of the alignment of the pointers in the latter case. It may also be necessary to insert dummy data at the end of each row to make the row length a multiple of four bytes.

Using __builtin_aligned

To avoid the need to use IPA to propagate alignment, and for situations when IPA cannot guarantee the alignment (but you can), it is possible to use the `__builtin_aligned` function to assert the alignment of important pointers, meaning that the pointer points to data that is aligned. It is important to remember when adding this declaration that you are responsible for making sure it is valid, and that, if the assertion is not true, the code produced by the compiler is likely to malfunction.

The assertion is particularly useful for function parameters, although you may assert that any pointer is aligned. For example, when compiling the function:

```
void copy(char *a, char *b) {
   int i;
   for (i=0; i<100; i++)
      a[i] = b[i];
}
```

Bad: Without IPA, compiler does not know the alignment of `a` and `b`.

the compiler does not know the alignment of pointers `a` and `b` if IPA is not being used. However, by modifying the function to:

```
void copy(char *a, char *b) {
   int i;
   __builtin_aligned(a, 4);
   __builtin_aligned(b, 4);
   for (i=0; i<100; i++)
```

```
        a[i] = b[i];
}
```

Good: Both pointer parameters are known to be aligned.

the compiler is told that the pointers are aligned on word boundaries. To assert instead that both pointers are always aligned one `char` past a word boundary, use:

```
void copy(char *a, char *b) {
int i;
__builtin_aligned(a+1, 4);
__builtin_aligned(b+1, 4);
for (i=0; i<100; i++)
   a[i] = b[i];
}
```

Good: Both pointer parameters are known to be misaligned.

The expression used as the first parameter to the built-in function obeys the usual C rules for pointer arithmetic. The second parameter should give the alignment in bytes as a literal constant.

Avoiding Aliases

It may seem that the iterations may be performed in any order in the following loop:

```
void fn(char a[], char b[], int n) {
int i;
for (i = 0; i < n; ++i)
   a[i] = b[i];
}
```

Bad: `a` and `b` may alias each other.

but `a` and `b` are both parameters, and, although they are declared with `[]`, they are in fact pointers, which may point to the same array. When the same data may be reachable through two pointers, they are said to alias each other.

If IPA is enabled, the compiler looks at the call sites of `fn` and tries to determine whether `a` and `b` can ever point to the same array.

Even with IPA, it is quite easy to create what appear to the compiler as aliases. The analysis works by associating pointers with sets of variables that they may refer to at some point in the program. If the sets for two pointers are found to intersect, then both pointers are assumed to point to the union of the two sets.

If `fn` above were called in two places with global arrays as arguments, then IPA would have the results shown below:

```
fn(glob1, glob2, N);
fn(glob1, glob2, N);
```

Good: sets for `a` and `b` do not intersect: `a` and `b` are not aliases.

```
fn(glob1, glob2, N);
fn(glob3, glob4, N);
```

Good: sets for `a` and `b` do not intersect: `a` and `b` are not aliases.

```
fn(glob1, glob2, N);
fn(glob3, glob1, N);
```

Bad: sets intersect - both `a` and `b` may access `glob1`; `a` and `b` may be aliases.

The third case arises because IPA considers the union of all calls at once, rather than considering each call individually, when determining whether there is a risk of aliasing. If each call were considered individually, IPA would have to take flow control into account and the number of permutations would make compilation time impracticably long.

The lack of control flow analysis can also create problems when a single pointer is used in multiple contexts. For example, it is better to write

```
int *p = a;
int *q = b;
   // some use of p
   // some use of q
```

Good: `p` and `q` do not alias.

than

```
int *p = a;
   // some use of p
p = b;
   // some use of p
```

Bad: Uses of `p` in different contexts may alias.

because the latter may cause extra apparent aliases between the two uses.

Indexed Arrays Versus Pointers

The C language allows a program to access data from an array in two ways: either by indexing from an invariant base pointer, or by incrementing a pointer. The following two versions of vector addition illustrate the two styles.

Style 1: using indexed arrays

```
void va_ind(const short a[], const short b[], short out[], int n) {
   int i;
   for (i = 0; i < n; ++i)
      out[i] = a[i] + b[i];
}
```

Style 2: using pointers

```
void va_ptr(const short a[], const short b[], short out[], int n) {
   int i;
```

```
    short *pout = out;
    const short *pa = a, *pb = b;
    for (i = 0; i < n; ++i)
       *pout++ = *pa++ + *pb++;
}
```

Trying Pointer and Indexed Styles

One might hope that the chosen style would not make any difference to the generated code, but this is not always the case. Sometimes, one version of an algorithm generates better optimized code than the other, but it is not always the same style that is better.

Try both pointer and index styles.

The pointer style introduces additional variables that compete with the surrounding code for resources during the compiler optimizer's analysis. Array accesses, on the other hand, must be transformed to pointers by the compiler—sometimes this is accomplished better by hand.

The best strategy is to start with array notation. If the generated code looks unsatisfactory, try using pointers. Outside the critical loops, use the indexed style, since it is easier to understand.

Function Inlining

Function inlining may be used in two ways:

- By annotating functions in the source code with the `inline` keyword. In this case, function inlining is only performed when optimization is enabled.
- By turning on automatic inlining with the `-Oa` switch (on page 1-46). This switch automatically enables optimization.

Inline small, frequently-executed functions.

You can use the compiler's `inline` keyword to indicate that functions should have code generated inline at the point of call. Doing this avoids various costs such as program flow latencies, function entry and exit instructions, and parameter passing overheads.

Using an `inline` function also has the advantage that the compiler can optimize through the inline code and does not have to assume that scratch registers and condition states are modified by the call. Prime candidates for inlining are small, frequently-used functions because they cause the least code-size increase while giving most performance benefit.

As an example of the usage of the `inline` keyword, the function below sums two input parameters and returns the result.

```
inline int add(int a, int b) {
   return (a+b);
}
```

Good: use of the `inline` keyword.

Inlining has a code size-to-performance trade-off that should be considered when it is used. With `-Oa`, the compiler automatically inlines small functions where possible. If the application has a tight upper code-size limit, the resulting code-size expansion may be too great. It is worth considering using automatic inlining in conjunction with the `-Ov n` switch (on page 1-47) to restrict inlining (and other optimizations with a code-size cost) to parts of the application that are performance-critical. This is considered in more detail later in this chapter.

Using Inline asm Statements

The compiler allows use of inline `asm` statements to insert small sections of assembly into C code.

Avoid use of inline `asm` statements where built-in functions may be used instead.

The compiler does not intensively optimize code that contains inline `asm` statements because it has little understanding about what the code in the statement does. In particular, use of an `asm` statement in a loop may inhibit useful transformations.

The compiler has been enhanced with a large number of built-in functions. These functions generate specific hardware instructions and are designed to allow the programmer to more finely tune the code produced by the compiler, or to allow access to system support functions. A complete list of compiler's built-in functions is given in "Compiler Built-In Functions" on page 1-103.

Use of these built-in functions is much preferred to using inline `asm` statements. Since the compiler knows what each built-in does, it can easily optimize around them. Conversely, since the compiler does not parse `asm` statements, it does not know what they do, and therefore is hindered in optimizing code that uses them. Note also that errors in the text string of an `asm` statement are caught by the assembler and not the compiler.

Examples of efficient use of built-in functions are given in "System Support Built-In Functions" on page 2-37.

Memory Usage

The compiler, in conjunction with the use of the linker description file (`.LDF`), allows the programmer control over data placement in memory. This section describes how to best lay out data for maximum performance.

Try to put arrays into different memory sections.

The processor hardware can support two memory operations on a single instruction line, combined with a compute instruction. Two memory operations will only complete in one cycle if the two addresses are situated in different memory blocks; if both access the same block, then a stall is incurred.

Take as an example the dot product loop below. Because data is loaded from both array `a` and array `b` in every iteration of the loop, it may be useful to ensure that these arrays are located in different blocks.

```
for (i=0; i<100; i++)
   sum += a[i] * b[i];
```

Bad: compiler assumes that two memory accesses together may give a stall.

First, define two memory banks in the `MEMORY` portion of the `.LDF` file.

Example: `MEMORY` portion of the `.LDF` file modified to define memory banks.

```
MEMORY {
   BANK_A1 {
      TYPE(RAM) WIDTH(8)
      START(start_address_1) END(end_address_1)
   }
   BANK_A2 {
      TYPE(RAM) WIDTH(8)
      START(start_address_2) END(end_address_2)
   }
}
```

Then, configure the `SECTIONS` portion to tell the linker to place data sections in specific memory banks.

Example: `SECTIONS` portion of the `.LDF` file modified to define memory banks.

```
SECTIONS {
   bank_a1 {
      INPUT_SECTION_ALIGN(4)
      INPUT_SECTIONS($OBJECTS(bank_a1))
   } >BANK_A1
   bank_a2 {
      INPUT_SECTION_ALIGN(4)
      INPUT_SECTIONS($OBJECTS(bank_a2))
   } >BANK_A2
}
```

In the C source code, declare arrays with the `section("section_name")` construct preceding a buffer declaration; in this case,

```
section("bank_a1") short a[100];
section("bank_a2") short b[100];
```

This ensures that two array accesses may occur simultaneously without incurring a stall.

Note that the explicit placement of data in sections can only be done for global data.

Using the Bank Qualifier

The bank qualifier can be used to write functions that use the fact that buffers are placed in separate memory blocks. For example, it might be useful to create a function,

```
void func(int bank("red") *p, int bank("blue") *q) {
     // some code
}
```

Good: uses bank qualifier to allow simultaneous access to `p` and `q`.

if you would like to call `func` in different places, but always with pointers to buffers in different sections of memory.

The bank qualifier symbolizes that the buffers are in different sections without requiring that the sections themselves be specified.

Therefore, `func` may be called with the first parameter pointing to memory in `section("bank_a1")` and the second pointing to data in `section("bank_a2")` or vice versa. Note that it is still necessary to explicitly place the data buffers in the memory sections; the bank qualifier merely informs the compiler that it may assume this has been done in order to generate more efficient code. Refer to "Bank Type Qualifiers" on page 1-96 for more information.

Improving Conditional Code

When compiling conditional statements, the compiler attempts to predict whether the condition will usually evaluate to true or to false, and will arrange for the most efficient path of execution to be that which is expected to be most commonly executed.

You can use the `expected_true` and `expected_false` built-in functions to control the compiler's behavior for specific cases. By using these functions, you can tell the compiler which way a condition is most likely to evaluate, and so influence the default flow of execution. For example,

```
if (buffer_valid(data_buffer))
   if (send_msg(data_buffer))
      system_failure();
```

shows two nested conditional statements. If it was known that, for this example, `buffer_valid()` would usually return true, but that `send_msg()` would rarely do so, the code could be written as

```
if (expected_true(buffer_valid(data_buffer)))
   if (expected_false(send_msg(data_buffer)))
      system_failure();
```

See "Compiler Performance Built-in Functions" on page 1-139 (on `expected_true` and `expected_false` functions) for more information.

The compiler can also determine the most commonly-executed branches automatically, using Profile-Guided Optimization. See "Optimization Control" on page 1-71 for more details.

Loop Guidelines

Loops are where an application ordinarily spends the majority of its time. It is therefore useful to look in detail at how to help the compiler to produce the most efficient loop code.

Keeping Loops Short

For best code efficiency, loops should be as small as possible. Large loop bodies are usually more complex and difficult to optimize. Additionally, they may require register data to be stored in memory. This causes a decrease in code density and execution performance.

Avoiding Unrolling Loops

Do not unroll loops yourself.

Not only does loop unrolling make the program harder to read but it also prevents optimization by complicating the code for the compiler.

```
void va1(const short a[], const short b[], short c[], int n)
{
   int i;
   for (i = 0; i < n; ++i) {
   c[i] = b[i] + a[i];
   }
}
```

Good: the compiler unrolls if it helps.

```
void va2(const short a[], const short b[], short c[], int n)
{
   short xa, xb, xc, ya, yb, yc;
   int i;
   for (i = 0; i < n; i+=2) {
      xb = b[i]; yb = b[i+1];
      xa = a[i]; ya = a[i+1];
```

```
        xc = xa + xb; yc = ya + yb;
        c[i] = xc; c[i+1] = yc;
    }
}
```

Bad: harder for the compiler to optimize.

Avoiding Loop-Carried Dependencies

A loop-carried dependency exists when computations in a given iteration of a loop cannot be completed without knowledge of the values calculated in earlier iterations. When a loop has such dependencies, the compiler cannot overlap loop iterations. Some dependencies are caused by scalar variables that are used before they are defined in a single iteration. However, an optimizer can reorder iterations in the presence of the class of scalar dependencies known as *reductions*. These are loops that reduce a vector of values to a scalar value using an associative and commutative operator—a common case being multiply and accumulate.

```
for (i = 0; i < n; ++i)
   x = a[i] - x;
```

Bad: loop-carried dependence is a scalar dependency.

```
for (i = 0; i < n; ++i)
   x += a[i] * b[i];
```

Good: loop-carried dependence is a reduction.

In the first case, the scalar dependency is the subtraction operation; the variable `x` is modified in a manner that gives different results if the iterations are performed out of order. In contrast, in the second case the properties of the addition operator stipulate that the compiler can perform the iterations in any order and still get the same result. Other examples of operators which are reductions are `bitwise` and/or, and `min/max`. In par-

ticular, the existence of loop-carried dependencies that are not reductions is one criterion that prevents the compiler from vectorizing a loop—that is, to execute more than one iteration concurrently.

Avoiding Loop Rotation by Hand

Do not rotate loops by hand.

Programmers are often tempted to "rotate" loops in DSP code by "hand" attempting to execute loads and stores from earlier or future iterations at the same time as computation from the current iteration. This technique introduces loop-carried dependencies that prevent the compiler from rearranging the code effectively. However, it is better to give the compiler a "normalized" version, and leave the rotation to the compiler.

For example,

```
int ss(short *a, short *b, int n) {
   int sum = 0;
   int i;
   for (i = 0; i < n; i++) {
      sum += a[i] + b[i];
   }
   return sum;
}
```

Good: is rotated by the compiler.

```
int ss(short *a, short *b, int n) {
   short ta, tb;
   int sum = 0;
   int i = 0;
   ta = a[i]; tb = b[i];
   for (i = 1; i < n; i++) {
      sum += ta + tb;
      ta = a[i]; tb = b[i];
   }
   sum += ta + tb;
```

```
    return sum;
}
```

Bad: rotated by hand—hard for the compiler to optimize.

By rotating the loop, the scalar variables `ta` and `tb` have been added, introducing loop-carried dependencies.

Avoiding Array Writes in Loops

Other dependencies can be caused by writes to array elements. In the following loop, the optimizer cannot determine whether the load from `a` reads a value defined on a previous iteration or one that will be overwritten in a subsequent iteration.

```
for (i = 0; i < n; ++i)
   a[i] = b[i] * a[c[i]];
```

Bad: has array dependency.

The optimizer can resolve access patterns where the addresses are expressions that vary by a fixed amount on each iteration. These are known as "induction variables".

```
for (i = 0; i < n; ++i)
   a[i+4] = b[i] * a[i];
```

Good: uses induction variables.

Inner Loops Versus Outer Loops

Inner loops should iterate more than outer loops.

The optimizer focuses on improving the performance of inner loops because this is where most programs spend the majority of their time. It is considered a good trade-off for an optimization to slow down the code before and after a loop if it is going to make the loop body run faster. Therefore, try to make sure that your algorithm also spends most of its

time in the inner loop; otherwise it may actually be made to run slower by optimization. If there are nested loops (where the outer loop runs many times and the inner loop runs a small number of times), it may be possible to rewrite the loops so that the outer loop has the fewer iterations.

Avoiding Conditional Code in Loops

If a loop contains conditional code, control-flow latencies may incur large penalties if the compiler has to generate conditional jumps within the loop. In some cases, the compiler is able to convert `IF-THEN-ELSE` and `?:` constructs into conditional instructions. In other cases, it is able to relocate the expression evaluation outside of the loop entirely. However, for important loops, linear code should be written where possible.

There are several tricks that can be used to remove the necessity for conditional code. For example, there is hardware support for `min` and `max`. The compiler usually succeeds in transforming conditional code equivalent to `min` or `max` into the single instruction. With particularly convoluted code the transformation may be missed, in which case it is better to use `min` or `max` in the source code.

The compiler does not perform the loop transformation that interchanges conditional code and loop structures. Instead of writing

```
for (i=0; i<100; i++) {
   if (mult_by_b)
      sum1 += a[i] * b[i];
   else
      sum1 += a[i] * c[i];
}
```

Bad: loop contains conditional code.

it is better to write

```
if (mult_by_b) {
   for (i=0; i<100; i++)
      sum1 += a[i] * b[i];
```

```
} else {
   for (i=0; i<100; i++)
      sum1 += a[i] * c[i];
}
```

Good: two simple loops can be optimized well.

if this is an important loop.

Avoiding Placing Function Calls in Loops

The compiler usually is unable to generate a hardware loop if the loop contains a function call due to the expense of saving and restoring the context of a hardware loop. In addition to obvious function calls, such as `printf()`, hardware loop generation can also be prevented by operations such as division, modulus, and some type coercions. These operations may require implicit calls to library functions. For more details, see "Data Types" on page 2-8.

Avoiding Non-Unit Strides

If you write a loop, such as

```
for (i=0; i<n; i+=3) {
      // some code
}
```

Bad: non-unit stride means division may be required.

then in order for the compiler to turn this into a hardware loop, it needs to work out the loop trip count. To do so, it must divide `n` by 3. The compiler decides that this is worthwhile as it speeds up the loop, but as discussed above, division is an expensive operation. Try to avoid creating loop control variables with strides of non-unit magnitude.

In addition, try to keep memory accesses in consecutive iterations of an inner loop contiguous. This is particularly applicable to multi-dimensional arrays. Therefore,

```
for (i=0; i<100; i++)
   for (j=0; j<100; j++)
      sum += a[i][j];
```

Good: memory accesses contiguous in inner loop.

is likely to be better than

```
for (i=0; i<100; i++)
   for (j=0; j<100; j++)
      sum += a[j][i];
```

Bad: loop cannot be unrolled to use wide loads.

as the former is more amenable to vectorization.

Loop Control

Use `int` types for loop control variables and array indices.

Use automatic variables for loop control and loop exit test.

For loop control variables and array indices, it is always better to use signed `int`s rather than any other integral type. The C standard requires various type promotions and standard conversions that complicate the code for the compiler optimizer. Frequently, the compiler is still able to deal with such code and create hardware loops and pointer induction variables. However, it does make it more difficult for the compiler to optimize and may occasionally result in under-optimized code.

The same advice goes for using automatic (local) variables for loop control. It is easy for a compiler to see that an automatic scalar, whose address is not taken, may be held in a register during a loop. But it is not as easy when the variable is a global or a function static.

Therefore, code such as

```
for (i=0; i<globvar; i++)
   a[i] = 10;
```

Bad: may need to reload `globvar` on every iteration.

may not create a hardware loop if the compiler cannot be sure that the write into the array `a` does not change the value of the global variable. The `globvar` variable must be reloaded each time around the loop before performing the exit test.

In this circumstance, the programmer can make the compiler's job easier by writing:

```
int upper_bound = globvar;
for (i=0; i<upper_bound; i++)
   a[i] = 10;
```

Good: easily becomes hardware loop.

Using the Restrict Qualifier

The `restrict` qualifier provides one way to help the compiler resolve pointer aliasing ambiguities. Accesses from distinct restricted pointers do not interfere with each other. The loads and stores in the following loop

```
for (i=0; i<100; i++)
   a[i] = b[i];
```

Bad: possible alias of arrays `a` and `b`.

may be disambiguated by writing

```
int * restrict p = a;
int * restrict q = b;
for (i=0; i<100; i++)
   *p++ = *q++;
```

Good: `restrict` qualifier tells compiler that memory accesses do not alias.

The `restrict` keyword is particularly useful on function parameters. The `restrict` keyword currently has no effect when used on variables defined in a scope nested within the outermost scope of a function.

Using the Const Qualifier

By default, the compiler assumes that the data referenced by a pointer to `const` type does not change. Therefore, another way to tell the compiler that the two arrays `a` and `b` do not overlap is to use the `const` keyword.

```
void copy(short *a, const short *b) {
   int i;
   for (i=0; i<100; i++)
      a[i] = b[i];
}
```

Good: pointers disambiguated via `const` qualifier.

The use of the `const` keyword in the above example has a similar effect as the `no_alias` pragma (see "#pragma no_alias" on page 2-50). In fact, the `const` implementation is better since it also allows the optimizer to use the fact that accesses via `a` and `b` are independent in other parts of the code, not just the inner loop.

In C, it is legal, though bad programming practice, to use casts to allow the data pointed to by pointers to `const` type to change. This should be avoided since, by default, the compiler generates code that assumes `const` data does not change. If you have a program that modifies `const` data through a pointer, you can generate standard-conforming code by using the compile-time flag `-const-read-write`.

Avoiding Long Latencies

All pipelined machines introduce stall cycles when you cannot execute the current instruction until a prior instruction has exited the pipeline. For example, the Blackfin processor stalls for three cycles on a table lookup; `a[b[i]]` takes three cycles more than you would expect.

If a stall is seen empirically, but it is not obvious to you exactly why it is occurring, a good way to learn about the cause is the **Pipeline Viewer**. This can be accessed through **Debug Windows -> Pipeline Viewer** in the VisualDSP++ 4.0 IDDE. By single-stepping through the program, you may see where the stall occurs. Note that the **Pipeline Viewer** is only available within a simulator session.

Using Built-In Functions in Code Optimization

Built-in functions, also known as compiler intrinsics, provide a method for the programmer to efficiently use low-level features of the processor hardware while programming in C. Although this section does not cover all the built-in functions available (for more information, refer to "Compiler Built-In Functions" on page 1-103), it presents some code examples where implementation choices are available to the programmer.

Fractional Data

Fractional data, represented as an integral type, can be manipulated in two ways: one way is the use of long promoted shifts and multiply constructs and the second is the use of compiler built-in functions. The built-in functions are recommended as they give you the most control over your data. Let us consider the fractional dot product algorithm. This may be written as:

```
long dot_product (short *a, short *b) {
   int i;
   long sum=0;
   for (i=0; i<100; i++) {
   /* this line is performance critical */
      sum += (((long)a[i]*b[i]) << 1);
   }
   return sum;
}
```

Bad: uses shifts to implement fractional multiplication.

This presents some problems to the optimizer. Normally, the code generated here would be a multiply, followed by a shift, followed by an accumulation. However, the processor hardware has a fractional multiply/accumulate instruction that performs all these tasks in one cycle.

In the example code, the compiler recognizes this idiom and replaces the multiply followed by shift with a fractional multiply. In more complicated cases, where perhaps the multiply is further separated from the shift, the compiler may not detect the possibility of using a fractional multiply.

Moreover, the transformation may in fact be invalid since it turns non-saturating integer operations into saturating fractional ones. Therefore, the results may change if the summation overflows. The transformation is enabled by default since it usually is what the programmer intended.

The recommended coding style is to use built-in functions. In the following example, a built-in function is used to multiply fractional 16-bit data.

```
#include <math.h>
fract32 dot_product(fract16 *a, fract16 *b) {
   int i;
   fract32 sum=0;
   for (i=0; i<100; i++) {
      /* this line is performance critical */
      sum += __builtin_mult_fr1x32(a[i],b[i]);
   }
   return sum;
}
```

Good: uses builtins to implement fractional multiplication.

Note that the `fract16` and `fract32` types used in the example above are merely `typedefs` to C integer types used by convention in standard include files. The compiler does not have any in-built knowledge of these types and treats them exactly as the integer types that they are `typedefed` to.

System Support Built-In Functions

Built-in functions are also provided to perform low-level system management, in particular for the manipulation of system registers (defined in `sysreg.h`). It is usually better to use these built-in functions rather than inline `asm` statements.

The built-in functions cause the compiler to generate efficient inline instructions and their use often results in better optimization of the surrounding code at the point where they are used. Using built-in functions also result in improved code readability. For more information on supported built-in functions, refer to "Compiler Built-In Functions" on page 1-103.

Examples of the two styles are:

```
unsigned int get_cycles(void) {
   unsigned int ret_val;
   asm("%0 = CYCLES;" : "=d" (ret_val) : : );
   return ret_val;
}
```

Bad: uses inline `asm` statement.

```
#include <ccblkfn.h>
#include <sysreg.h>
unsigned int get_cycles(void) {
   return sysreg_read(reg_CYCLES);
}
```

Good: uses `sysreg.h`.

This example reads and returns the `CYCLES` register.

Using Circular Buffers

Circular buffers are often extremely useful in DSP-style code. They can be used in several ways. Consider the C code:

```
for (i=0; i<1000; i++) {
   sum += a[i] * b[i%20];
}
```

Good: the compiler knows that `b` is accessed as a circular buffer.

Clearly the access to array `b` is a circular buffer. When optimization is enabled, the compiler produces a hardware circular buffer instruction for this access.

Consider the slightly more complicated example:

```
for (i=0; i<1000; i+=n) {
   sum += a[i] * b[i%20];
}
```

Bad: may not be able to use circular buffer to access `b`.

In this case, the compiler does not know if `n` is positive and less than 20. If it is, then the access may be correctly implemented as a hardware circular buffer. On the other hand, if it is greater than 20, a circular buffer increment may not yield the same results as the C code.

The programmer has two options here. One is to compile with the `-force-circbuf` switch (on page 1-32). This tells the compiler that any access of the form `a[i%n]` should be considered as a circular buffer. Before using this switch, check that this assumption is valid for your application.

The preferred option is to use built-in functions to perform the circular buffering. Two functions (`__builtin_circindex` and `__builtin_circptr`) are provided for this purpose.

To make it clear to the compiler that a circular buffer should be used, you may write either:

```
for (i=0, j=0; i<1000; i+=n) {
   sum += a[i] * b[j];
   j = __builtin_circindex(j, n, 20);
}
```

Good: explicit use of circular buffer via `__builtin_circindex`.

or

```
int *p = b;
for (i=0, j=0; i<1000; i+=n) {
   sum += a[i] * (*p);
   p = __builtin_circptr(p, n, b, 80);
}
```

Good: explicit use of circular buffer via `__builtin_circptr`.

For more information, refer to "Circular Buffer Built-In Functions" on page 1-134).

Smaller Applications: Optimizing for Code Size

The same ethos for producing fast code also applies to producing small code. You should present the algorithm in a way that gives the optimizer excellent visibility of the operations and data, and hence the greatest freedom to safely manipulate the code to produce small applications.

Once the program is presented in this way, the optimization strategy depends on the code size constraint that the program must obey. The first step should be to optimize the application for full performance, using `-O` or `-ipa` switches. If this obeys the code size constraints, then no more need be done.

The "optimize for space" switch `-Os` (on page 1-47), which may be used in conjunction with IPA, performs every performance-enhancing transformation except those that increase code size. In addition, the `-e` linker switch (`-flags-link -e` if used from the compiler command line) may be helpful (on page 1-31). This operation performs section elimination in the linker to remove unneeded data and code. If the code produced with `-Os` and `-e` does not meet the code size constraint, some analysis of the source code is required to try to further reduce the code size.

Note that loop transformations such as unrolling and software pipelining increase code size. But it is these loop transformations that also give the greatest performance benefit. Therefore, in many cases compiling for minimum code size produces significantly slower code than optimizing for speed.

The compiler provides a way to balance between the two extremes of `-O` and `-Os`. This is the sliding-scale `-Ov num` switch (adjustable using the optimization slider bar under **Project Options** in the VisualDSP++ IDDE), described on page 1-47. The `num` parameter may be a value between 0 and 100, where the lower value corresponds to minimum code size and the upper to maximum performance. A value in-between is used

to optimize the frequently-executed regions of code for maximum performance, while keeping the infrequently-executed parts as small as possible. The switch is most reliable when using profile-guided optimization (see "Optimization Control" on page 1-71) since the execution counts of the various code regions have been measured experimentally. Without PGO, the execution counts are estimated, based on the depth of loop nesting.

Avoid the use of inline code.

Avoid using the `inline` keyword to inline code for functions that are used a number of times, especially if they not very small functions. The `-Os` switch does not have any effect on the use of the `inline` keyword. It does, however, prevent automatic inlining (using the `-Oa` switch) from increasing the code size. Macro functions can also cause code expansion and should be used with care.

See "Bit Fields" on page 2-11 for information on how bit fields affect code size.

Using Pragmas for Optimization

Pragmas can assist optimization by allowing the programmer to make assertions or suggestions to the compiler. This section looks at how they can be used to finely tune source code. Refer to "Pragmas" on page 1-148 for full details of how each pragma works. The emphasis of this section is to consider under what circumstances they are useful during the optimization process.

In most cases, the pragmas serve to give the compiler information which it is unable to deduce for itself. It must be emphasized that the programmer is responsible for making sure that the information given by the pragma is valid in the context in which it is used. Use of a pragma to assert that a function or loop has a quality that it does not in fact have is likely to result in incorrect code and hence a malfunctioning application.

An advantage of the use of pragmas is that they allow code to remain portable, since they normally are ignored by a compiler that does not recognize them.

Function Pragmas

Function pragmas include `#pragma alloc`, `#pragma const`, `#pragma pure`, `#pragma result_alignment`, `#pragma regs_clobbered`, and `#pragma optimize_{off|for_speed|for_space|as_cmd_line}`.

#pragma alloc

This pragma asserts that the function behaves like the `malloc` library function. In particular, it returns a pointer to new memory that cannot alias any pre-existing buffers. In the following code,

```
#pragma alloc
int *new_buf(void);
int *vmul(int *a, int *b) {
   int i;
```

```
    int *out = new_buf();
    for (i=0; i<100; i++)
       out[i] = a[i] * b[i];
}
```

Good: uses `#pragma alloc` to disambiguate out from `a` and `b`.

the use of the pragma allows the compiler to be sure that the write into buffer `out` does not modify either of the two input buffers `a` or `b`. Therefore, the iterations of the loop may be reordered.

#pragma const

This pragma asserts to the compiler that a function does not have any side effects (such as modifying global variables or data buffers), and the result returned is only a function of the parameter values. The pragma may be applied to a function prototype or definition. It helps the compiler since two calls to the function with identical parameters always yields the same result. In this way, calls to `#pragma const` functions may be hoisted out of loops if their parameters are loop independent.

#pragma pure

Like `#pragma const`, this pragma asserts to the compiler that a function does not have any side effects (such as modifying global variables or data buffers). However, the result returned may be a function of both the parameter values and any global variables. The pragma may be applied to a function prototype or definition. Two calls to the function with identical parameters always yields the same result provided that no global variables have been modified between the calls. Hence, calls to `#pragma pure` functions may be hoisted out of loops if their parameters are loop independent and no global variables are modified in the loop.

#pragma result_alignment

This pragma may be used on functions that have either pointer or integer results. When a function returns a pointer, the pragma is used to assert that the return result always has some specified alignment. Therefore, the example might be further refined if it is known that the `new_buf` function always returns buffers which are aligned on a word boundary.

```
#pragma alloc
#pragma result_alignment (4)
int *new_buf(void);

int *vmul(int *a, int *b) {
   int i;
   int *out = new_buf();
   for (i=0; i<100; i++)
      out[i] = a[i] * b[i];
}
```

Good: uses pragma `result_alignment` to specify that out has strict alignment.

Further details on this pragma may be found in "#pragma result_alignment (n)" on page 1-173. Another, more laborious way to achieve the same effect would be to use `__builtin_aligned` at every call site to assert the alignment of the returned result.

#pragma regs_clobbered

This pragma is a useful way to improve the performance of code that makes function calls. The best use of the pragma is to increase the number of call-preserved registers available across a function call. There are two complementary ways in which this may be done.

First of all, suppose you have a function written in assembly that you wish to call from C source code. The `regs_clobbered` pragma may be applied to the function prototype to specify which registers are "clobbered" by the

assembly function, that is, which registers may have different values before and after the function call. Consider for example a simple assembly function that adds two integers and then masks the result to fit into 8 bits:

```
_add_mask:
   R0 = R0 + R1;
   R0 = R0.B (z);
   RTS;
._add_mask.end
```

Clearly the function does not modify the majority of the scratch registers available and thus these could instead be used as call-preserved registers. In this way, fewer spills to the stack are needed in the caller function. Using the following prototype,

```
#pragma regs_clobbered "R0, ASTAT"
int add_mask(int, int);
```

Good: uses `regs_clobbered` to increase call-preserved register set.

the compiler is told which registers are modified by a call to the `add_mask` function. The registers not specified by the pragma are assumed to preserve their values across such a call and the compiler may use these spare registers to its advantage when optimizing the call sites.

The pragma is also powerful when all of the source code is written in C. In the above example, a C implementation might be:

```
int add_mask(int a, int b) {
   return ((a+b)&255);
}
```

Bad: function thought to clobber entire volatile register set.

Since this function does not need many registers when compiled, it can be defined using

```
#pragma regs_clobbered "R0, CCset"
int add_mask(int a, int b) {
```

```
    return ((a+b)&255);
}
```

Good: function compiled to preserve most registers.

to ensure that any other registers aside from `R0` and the condition codes are not modified by the function. If any other registers are used in the compilation of the function, they are saved and restored during the function prologue and epilogue.

In general, it is not very helpful to specify any of the condition codes as call-preserved as they are difficult to save and restore and are usually clobbered by any function. Moreover, it is usually of limited benefit to keep them live across a function call. Therefore, it is better to use `CCset` (all condition codes) rather than `ASTAT` in the clobbered set above. For more information, refer to "#pragma regs_clobbered string" on page 1-169.

#pragma optimize_{off|for_speed|for_space|as_cmd_line}

This pragma (pragmas) may be used to change the optimization setting on a function-by-function basis. In particular, it may be useful to optimize functions that are rarely called (for example, error handling code) for space (using `#pragma optimize_for_space`), whereas functions critical to performance should be compiled for maximum speed (`#pragma optimize_for_speed`). The `#pragma optimize_off` is useful for debugging specific functions without increasing the size or decreasing the performance of the overall application unnecessarily.

The `#pragma optimize_as_cmd_line` resets the optimization settings to those specified on the `ccblkfn` command line when the compiler is invoked. Refer to "General Optimization Pragmas" on page 1-159 for more information.

Loop Optimization Pragmas

Many pragmas are targeted towards helping to produce optimal code for inner loops. These are the `loop_count`, `no_vectorization`, `vector_for`, `all_aligned`, `different_banks`, and `no_alias` pragmas.

#pragma loop_count

This pragma enables the programmer to inform the compiler about a loop's iteration count. The compiler is able to make more reliable decisions about the optimization strategy for a loop if it knows the iteration count range. If you know that the loop count is always a multiple of some constant, this can also be useful as it allows a loop to be partially unrolled or vectorized without the need for conditionally-executed iterations. Knowledge of the minimum trip count may allow the compiler to omit the guards that are usually required after software pipelining. Any of the parameters of the pragma that are unknown may be left blank.

An example of the use of the `loop_count` pragma might be:

```
#pragma loop_count(/*minimum*/ 40, /*maximum*/ 100, /*modulo*/ 4)
for (i=0; i<n; i++)
   a[i] = b[i];
```

Good: the `loop_count` pragma gives the compiler helpful information to assist optimization.

For more information, refer to "#pragma loop_count(min, max, modulo)" on page 1-157.

#pragma no_vectorization

Vectorization (executing more than one iteration of a loop in parallel) can slow down loops with very small iteration counts since a loop prologue and epilogue are required. The `no_vectorization` pragma can be used directly above a `for or do` loop to instruct the compiler not to vectorize the loop.

#pragma vector_for

This pragma is used to help the compiler resolve dependencies that would normally prevent it from vectorizing a loop. It tells the compiler that all iterations of the loop may be run in parallel with each other, subject to rearrangement of reduction expressions in the loop. In other words, there are no loop-carried dependencies except reductions. An optional parameter, `n`, may be given in parentheses to say that only `n` iterations of the loop may be run in parallel. The parameter must be a literal value.

For example,

```
for (i=0; i<100; i++)
   a[i] = b[i] + a[i-4];
```

Bad: cannot be vectorized due to possible alias between `a` and `b`.

cannot be vectorized if the compiler cannot tell that array `b` does not alias array `a`. But the pragma may be added to instruct the compiler that in this case four iterations may be executed concurrently.

```
#pragma vector_for (4)
for (i=0; i<100; i++)
   a[i] = b[i] + a[i-4];
```

Good: pragma `vector_for` disambiguates alias.

Note that this pragma does not force the compiler to vectorize the loop; the optimizer checks various properties of the loop and does not vectorize it if it believes that it is unsafe or cannot deduce information necessary to carry out the vectorization transformation. The pragma assures the compiler that there are no loop-carried dependencies, but there may be other properties of the loop that prevent vectorization.

In cases where vectorization is impossible, the information given in the assertion made by `vector_for` may still be put to good use in aiding other optimizations.

For more information, refer to “#pragma vector_for” on page 1-159.

#pragma all_aligned

This pragma is used as shorthand for multiple `__builtin_aligned` assertions. By prefixing a `for` loop with the pragma, it is asserted that every pointer variable in the loop is aligned on a word boundary at the beginning of the first iteration. Therefore, adding the pragma to the following loop

```
#pragma all_aligned
for (i=0; i<100; i++)
   a[i] = b[i];
```

Good: uses `all_aligned` to inform compiler of alignment of `a` and `b`.

is equivalent to writing

```
__builtin_aligned(a, 4);
__builtin_aligned(b, 4);
for (i=0; i<100; i++)
   a[i] = b[i];
```

Good: uses `__builtin_aligned` to give alignment of `a` and `b`.

In addition, the `all_aligned` pragma may take an optional literal integer argument `n` in parentheses. This tells the compiler that all pointer variables are aligned on a word boundary at the beginning of the n^{th} iteration. Note that the iteration count begins at zero.

Therefore,

```
#pragma all_aligned (3)
for (i=99; i>=0; i--)
   a[i] = b[i];
```

Good: uses `all_aligned` to inform compiler of alignment of `a` and `b`.

is equivalent to

```
__builtin_aligned(a+96, 4);
__builtin_aligned(b+96, 4);
```

```
for (i=99; i>=0; i--)
   a[i] = b[i];
```

Good: uses `__builtin_aligned` to give alignment of `a` and `b`.

For more information, refer to "Using __builtin_aligned" on page 2-15.

#pragma different_banks

This pragma is used as shorthand for declaring multiple pointer types with different bank qualifiers. It asserts that any two independent memory accesses in the loop may be issued together without incurring a stall. Therefore, writing

```
#pragma different_banks
for (i=0; i<100; i++)
   a[i] = b[i];
```

Good: uses different banks to allow simultaneous accesses to `a` and `b`.

allows a single instruction loop to be created if it is known that `a` and `b` do not alias each other. See "#pragma different_banks" on page 1-157 for more information.

#pragma no_alias

When immediately preceding a loop, the `no_alias` pragma asserts that no load or store in the loop accesses the same memory. This helps to produce shorter loop kernels as it permits instructions in the loop to be rearranged more freely. See "#pragma no_alias" on page 1-158 for more information.

Useful Optimization Switches

Table 2-1 lists the compiler switches useful during the optimization process.

Table 2-1. C/C++ Compiler Optimization Switches

Switch Name	Description
`-const-read-write` on page 1-28	Specifies that data accessed via a pointer to `const` data may be modified elsewhere
`-flags-link -e` on page 1-31	Specifies linker section elimination
`-force-circbuf` on page 1-32	Treats array references of the form `array[i%n]` as circular buffer operations
`-ipa` on page 1-36	Turns on inter-procedural optimization. Implies use of `-O`. May be used in conjunction with `-Os` or `-Ov`.
`-no-fp-associative` on page 1-43	Does not treat floating-point multiply and addition as an associative
`-O` on page 1-45	Enables code optimizations and optimizes the file for speed
`-Os` on page 1-47	Optimizes the file for size
`-Ov num` on page 1-47	Controls speed vs. size optimizations (sliding scale)
`-pguide` on page 1-50	Adds instrumentation for the gathering of a profile as the first stage of performing profile-guided optimization
`-save-temps` on page 1-54	Saves intermediate files (for example, `.s`)

Example: How the Optimizer Works

The following fractional scalar product loop is used to show how the optimizer works.

Example: C source code for fixed-point scalar product.

```
#include <fract.h>
fract32 sp(fract16 *a, fract16 *b) {
   int i;
   fract32 sum=0;
   __builtin_aligned(a, 4);
   __builtin_aligned(b, 4);
   for (i=0; i<100; i++) {
      sum = __builtin_add_fr1x32(sum,
            __builtin_mult_fr1x32(a[i],b[i]));
   }
   return sum;
}
```

After code generation and conventional scalar optimizations are done, the compiler generates a loop that looks something like the following example.

Example: Initial code generated for fixed-point scalar product.

```
    P2 = 100;
    LSETUP(.P1L3, .P1L4 - 2) LC0 = P2;
.P1L3:
    R0 = W[P0++] (X);
    R2 = W[P1++] (X);
    A0 += R0.L * R2.L;
.P1L4:
   R0 = A0.w;
```

The loop exit test has been moved to the bottom and the loop counter rewritten to count down to zero, allowing a zero-overhead loop to be generated. The `sum` is being accumulated in `A0`. The `P0` and `P1` hold pointers that are initialized with the parameters `a` and `b`, respectively, and are incremented on each iteration.

In order to use 32-bit memory accesses, the optimizer unrolls the loop to run two iterations in parallel. The `sum` is now being accumulated in `A0` and `A1`, which must be added together after the loop to produce the final result. In order to use word loads, the compiler has to know that `P0` and `P1` have initial values that are multiples of four bytes.

This is done in the example by use of `__builtin_aligned`, although it could also have been propagated with IPA. Note also, that unless the compiler knows that the original loop was executed an even number of times, a conditionally-executed odd iteration must be inserted outside the loop.

Example: Code generated for fixed-point scalar product after vectorization transformation.

```
      P2 = 50;
      A1 = A0 = 0;
      LSETUP(.P1L3, .P1L4 - 4) LC0 = P2;
   .P1L3:
      R0 = [P0++];
      R2 = [P1++];
      A1+=R0.H*R2.H, A0+=R0.L*R2.L;
   .P1L4:
      R0 = (A0+=A1);
```

Finally, the optimizer rotates the loop, unrolling and overlapping iterations to obtain the highest possible use of functional units. Code similar to the following is generated.

Example: Code generated for fixed-point scalar product after software pipelining.

```
      A1=A0=0 || R0 = [P0++] || NOP;
      R2 = [I1++];
      P2 = 49;
      LSETUP(.P1L3,.P1L4-8) LC0 = P2;
   .P1L3:
      A1+=R0.H*R2.H, A0+=R0.L*R2.L || R0 = [P0++] || R2 = [I1++];
   .P1L4:
      A1+=R0.H*R2.H, A0+=R0.L*R2.L;
      R0 = (A0+=A1);
```

Assembly Optimizer Annotations

When the compiler optimizations are enabled the compiler can perform a large number of optimizations in order to generate the resultant assembly code. The decisions taken by the compiler as to whether certain optimizations are safe or worthwhile are generally invisible to a programmer, though it could be of benefit to programmers to be given feedback from the compiler with regards to the decisions made during the optimization phase. The intention of the information provided is to give a programmer an understanding of how close to optimal a program is and what more could possibly be done to improve the generated code.

The feedback from the compiler optimizer is provided by means of annotations made to the assembly file generated by the compiler. The assembly file generated by the compiler can be kept by specifying the `-S` switch (see on page 1-53), the `-save-temps` switch (see on page 1-54) or by checking the **Project Options->Compile->General->Save temporary files** option in VisualDSP++ IDDE.

There are several areas of information provided by the assembly annotations that could be of benefit in code evaluation to improve the generated code, for example providing an indication of resource usage or the absence of a particular optimization from the resultant code, which can often be more important than its presence. The intention of the assembly code annotations is to give the programmer some insight as to the reasons why the compiler enables and disables certain optimizations for a specific code sequence.

The assembly code generated by the compiler optimizer is annotated with the following information:

- “Procedure Statistics” on page 2-55
- “Loop Identification” on page 2-58
- “Vectorization” on page 2-65

The assembly output for the examples in this chapter may differ based on optimization flags and the version of the compiler. As a result, you may not be able to exactly reproduce these results.

Procedure Statistics

For each function call, the following is reported:

- Frame size – size of stack frame.
- Registers used – Since function calls tend to implicitly clobber registers, there are several sets:

    1. The first set is composed of the scratch registers changed by the current function, not counting the registers that are implicitly clobbered by the functions called from the current function.
    2. The second set are the call preserved registers changed by the current function, not counting the registers that are implicitly clobbered by the functions called from the current function.
    3. The third set are the registers clobbered by the inner function calls.

Example A

Consider the following program:

```
struct str {
    int x1, x2;
};

int func1(struct str*, int *);
int func2(struct str s);
int foo(int in)
{
```

```
    int sum = 0;
    int local;
    struct str l_str;
    sum += func1(&l_str, &local);
    sum += func2(l_str);
    return sum;
}
```

The procedure statistics for `foo` are:

```
_foo:
//-------------------------------------------------------------
//.................Procedure statistics:
//-------------------------------------------------------------
//
//   Frame size            = 36 words
//
// Scratch registers modified:{R0-R1,P0,ASTAT}
//
// Call preserved registers used:{R7}
//
// Registers clobbered by function calls:
// {R0-R3,P0-P2,I0-I3,B0-B3,M0-M3,ASTAT,SEQSTAT,RETS,CC,
// A0-A1,LC0-LC1,LT0-LT1,LB0-LB1,SYSCFG,CYCLES,CYCLES2}
//-------------------------------------------------------------
// "foo.c" line 9 col 1
           link  28;
//" foo.c" line 14 col 17
//         -- 3 bubbles --;
           R0 = FP;
           R0 +=  -4;
           R1 = ROT R0 by 0  || [SP+ 12] = R7 || NOP;
           [SP+ 4] = R0;
           R0 +=  -8;
           CALL.X _func1;
           R7 = ROT R0 by 0  || R0 = [FP+ -12] || NOP;
//" foo.c" line 15 col 17
           R1 = [FP+ -8];
           CALL.X _func2;
//" foo.c" line 15 col 5
```

```
            R0 = R7 + R0 (NS) || P0 = [FP+ 4] || NOP;
//" foo.c" line 17 col 1
            R7 = [SP+ 12];
            unlink;
//          -- 2 bubbles --;
            JUMP (P0);

._foo.end:
      .global _foo;
      .type _foo,STT_FUNC;
```

Notes:

- The Frame size is 36 bytes, indicating how much space is allocated on the stack by the function. The frame size includes:
  - 4 bytes for RETS
  - 4 bytes for the Frame Pointer
  - the space allocated by the compiler, for local variables (12 bytes: 4 for local, 8 for `l_str`);
  - space required to save any callee-preserved registers (4 bytes, for `R7`)
  - space required for parameters being passed to functions called by this one (12 bytes)
- The scratch registers modified directly by the function. These are the registers the compiler does not need to save, before modifying. In this case, the registers are `R0`, `R1`, `P0` and `ASTAT`. This does not include any registers that are only modified by calls to other functions.

- The callee-preserved registers used by the function. These registers must be saved before modification, and restored afterwards. In this case, the compiler uses `R7`, and the saved value for `R7` accounts for 4 bytes of frame size.
- The registers clobbered by function calls made by this function. This is the union of all the registers that will be modified by the calls to other functions. In this case, the registers are the default scratch register set, modified by calls to `func1` and `func2`.

Loop Identification

One very useful annotation is loop identification—that is, showing the relationship between the source program loops and the generated assembly code. This is not easy due to the various loop optimizations. Some of the original loops may not be present, being entirely unrolled or merged with other loops. It can be extremely difficult describing what has happened to them.

Finally, the assembly code may contain compiler-generated loops that do not correspond to any loop in the user program, but rather represent constructs such as structure assignment or calls to `memcpy`.

Loop Identification Annotations

Loop identification annotation rules are:

- Annotate only the loops that originate from the C looping constructs `do`, `while`, and `for`. Therefore, any `goto` defined loop is not accounted for.
- A loop is identified by the position of the corresponding keyword (`do`, `while`, `for`) in the source file.
- Account for all such loops in the original user program.

- Generally, loop bodies are delimited between the `Lx: Loop` at `<file position>` and `End Loop Lx` assembly annotation. The former annotation follows the label of the first block in the loop. The later annotation follows the first jump back to the beginning of the loop. However, there are cases when the code corresponding to a user loop cannot be entirely represented between such two markers. In such cases the assembly code contains blocks that belong to a loop, but are not contained between that loop's end markers. Such blocks are annotated with a comment identifying the innermost loop they belong to `Part of Loop Lx`.
- Sometimes a loop in the original program does not show up in the assembly file because it was either transformed or deleted. In either case, a short description of what happened to the loop is given at the beginning of the function.
- A program's innermost loops are those loops that do not contain other loops. In addition to regular loop information, the innermost loops with no control flow and no function calls are annotated with additional information such as:
    - Cycle count: the number of cycles needed to execute one iteration of the loop, including the stalls.
    - Resource usage: the resources used during one iteration of the loop. For each resource we show how many of that resource are used, how many are available and the percentage of utilization during the entire loop. Resources are shown in decreasing order of utilization. Note that 100% utilization means that the corresponding resource is used at its full capacity and represents a bottleneck for the loop.
- Some loops are subject to optimizations such as vectorization. These loops receive additional annotations as described in the vectorization section.

Resource Definitions

For each cycle, a Blackfin processor may execute a single 16- or 32-bit instruction, or it may execute a 64-bit multi-issued instruction, consisting of a 32-bit instruction and two 16-bit instructions. In either case, at most one store instruction may be executed. Not all 16-bit instructions are valid for the multi-issue slots, and not all of those may be placed into either slot. Consequently, the resources are divided into Group 1 (use of the first 16-bit multi-issue slot), and Group 1 or 2 (use of either 16-bit multi-issue slot).

The resource usage is described in terms of missed opportunities by the compiler, that is, slots where the compiler has had to issue a `NOP` or `MNOP` instruction.

An instruction of the form:

```
R0 = R0 + R1 (NS) || R1 = [P0++] || NOP;
```

has managed to use both the 32-bit ALU slot and one of the 16-bit memory access slots, but has not managed to use the second 16-bit memory access slot. Therefore, this counts as:

- 1 out of 1 possible 32-bit ALU/MAC instructions
- 1 out of 1 possible Group 1 instructions
- 1 out of 2 possible Group 1 or 2 instructions
- 0 out of 1 possible stores

A single-issued instruction is seen as occupying all issue-slots at once, because the processor cannot issue other instructions in parallel. Consequently, there are no opportunities missed by the compiler. Thus, a single-issue instruction such as:

```
R2 = R0 + R1 ;
```

is counted as:

- 1 out of 1 possible 32-bit ALU/MAC instructions
- 1 out of 1 possible Group 1 instructions
- 2 out of 2 possible Group 1 or 2 instructions
- 1 out of 1 possible stores

in that the compiler has not had to issue NOPs or MNOPs, and so no resources have been unutilised.

Example B

Consider the following example:

```
1   int bar(int a[10000])
2   {
3      int i, sum = 0;
4      for (i = 0; i < 9999; ++i)
5          sum += (sum + 1);
6      while (i-- < 9999) /* this loop doesn't get executed */
7          a[i] = 2*i;
8      return sum;
9 }
```

The two loops are accounted for as follows:

```
_bar:
//--------------------------------------------------------------
.................. procedure statistics .........
//--------------------------------------------------------------
// Original Loop at " exampleC.c" line 6 col 6 -- loop structure
   removed due to constant propagation.
//--------------------------------------------------------------
//" exampleC.c" line 2 col 1
//" exampleC.c" line 4 col 6
                P0 = 9999 (X);
//" exampleC.c" line 2 col 1
```

```
                R0 =  0;
                R1 =  1;
//" exampleC.c" line 4 col 6
                LSETUP (._P1L2 , ._P1L3-2) LC0 = P0;

.P1L2:
//-------------------------------------------------------------
// Loop at " exampleC.c" line 4 col 6
//-------------------------------------------------------------
// This loop executes 1 iteration of the original loop in 2 cycles.
//-------------------------------------------------------------
// This loop's resource usage is:
//    32-bit ALU/MAC used 2 out of 2 (100.0%)
//    Group 1 used 2 out of 2 (100.0%)
//    Group 1 or 2 used 4 out of 4 (100.0%)
//    Store used 2 out of 2 (100.0%)
//-------------------------------------------------------------
                R2 = R0 + R1;
//" exampleC.c" line 5 col 10
                R0 = R0 + R2;
//" exampleC.c" line 4 col 6
//      end loop ._P1L2;
//-------------------------------------------------------------
// End Loop L2

.P1L3:
//-------------------------------------------------------------
// Part of top level (no loop)
//-------------------------------------------------------------
//" exampleC.c" line 9 col 1
                RTS;

    _bar.end:
```

Notes:

- The keywords identifying the two loops are:
  - `for` – whose position is in the file `exampleC.c`, line 4, column 5
  - `while` – whose position is in the file `exampleC.c`, line 6, column 6
- Immediately after the procedure statistics, there is a message stating that the loop at line 6 in the user program was removed. The reason was constant propagation, which in this case realizes that the value of `i` after the first loop is 9999 and the second loop does not get executed.
- The start of the loop at line 4 is marked in the assembly by the "Loop at exampleC.c, line 4, column 5" annotation. This annotation follows the loop label `.P1L2`. The loop label `End Loop L2` is used to identify the end of the loop.
- The loop resource information accounts for all instructions and stalls inside the loop. In this particular case, the loop body is executed in two cycles, one instruction for each cycle. Both instructions are single-issue instructions, and the compiler has not found it necessary to issue any NOPs or MNOPs, so it reports full utilitization.

File Position

As seen in the Example C, a file position is given using the file name, line number and the column number in that file: `"exampleC.c" line 4 col 5`.

This scheme uniquely identifies a source code position, unless inlining is involved. In presence of inlining, a piece of code from a certain file position can be inlined at several places, which in turn can be inlined at other places. Since inlining can happen for an unspecified number of times, a recursive scheme is used to describe a general file position.

Therefore, a <general file position> is <file position> inlined from <general file position>.

Example C (Inlining)

Consider the following source code:

```
5   void f2(int n);
6   inline void f3(int n)
7   {
8      while(n--)
9           f4();
10        if (n == 7)
11           f2(3*n);
12  }
13
14  inline void f2(int n)
15  {
16     while(n--) {
17          f3(n);
18          f3(2*n);
19     }
20  }
21  void f1(volatile unsigned int i)
22  {
23     f2(30);
24  }
```

Here is some of the code generated for function f1:

```
//------------------------------------------------------------
// Part of Loop 1, depth 1
//------------------------------------------------------------
// "exampleD.c" line 11 col 11
                R0 = R4 <<  1;
                R0 = R0 + R4;
// "exampleD.c" line 11 col 13
                CALL.X _f2;
```

```
                JUMP ._P2L11;

._P2L7:
//-------------------------------------------------------------
// "exampleD.c" line 8 col 4
                R6 = ROT R7 by 0  || R5 = [FP+ -8] || NOP;

._P2L9::
//-------------------------------------------------------------
// Loop at " exampleD.c" line 8 col 4 inlined from "exampleD.c"
// line 18 col 11 inlined from "exampleD.c" line 23 col 6
//-------------------------------------------------------------

//------------------------------------ --------------------
// End Loop L9
//-------------------------------------------------------------
//------------------------------------ --------------------
// Part of Loop 1, depth 1
//------------------------------------ --------------------
JUMP ._P2L8;
```

Vectorization

The trip count of a loop is the number of times the loop body gets executed.

Under certain conditions, the compiler is able to take two operations executed in consecutive iterations of a loop and to execute them in only one more powerful instruction giving a loop with a smaller trip count. The transformation in which operations from two subsequent iterations are executed in one more powerful single operation is called vectorization.

For instance, the original loop may start with a trip count of 1000.

```
for(i=0; i< 1000; ++i)
   a[i] = b[i] + c[i];
```

and, after the optimization, end up with the vectorized loop with a final trip count of 500. The vectorization factor is the number of operations in the original loop that are executed at once in the transformed loop. It is illustrated using some pseudo code below.

```
for(i=0; i< 500; i+=2)
   (a[i], a[i+1]) = (b[i],b[i+1]) .plus2. (c[i], c[i+1]);
```

In the above example, the vectorization factor is 2. A loop may be vectorized more than once.

If the trip count is not a multiple of the vectorization factor, some iterations need to be peeled off and executed unvectorized. Thus, if in the previous example, the trip count of the original loop was 1001, then the vectorized code would be:

```
for(i=0; i< 500; i+=2)
   (a[i], a[i+1]) = (b[i],b[i+1]) .plus2. (c[i], c[i+1]);
a[1000] = b[1000] + c[1000];
// This is one iteration peeled from
// the back of the loop.
```

In the above examples, the trip count is known and the amount of peeling is also known. If the trip count is not known (it is a variable), the number of peeled iterations depends on the trip count, and in such cases, the optimized code contains peeled iterations that are executed conditionally.

Unroll and Jam

Another vectorization related transformation is unroll and jam. Let's consider the following function:

```
/* unroll and jam example */
void f_unroll_and_jam(short a[][40], short *restrict c)  {
   __builtin_aligned(a, 4);
   __builtin_aligned(c, 4);
   for (i=0; i<60; i++) {
       short sum=0;
       for (j=0; j<40; j++) {
```

```
            sum += a[j][i];
        }
        c[i] = sum;
    }
}
```

The outer loop can be unrolled twice and the result is:

```
void f_unroll_and_jam(int a[][40], short *restrict c) {
   int i, j;

   __builtin_aligned(a, 4);
   __builtin_aligned(c, 4);
   for (i=0; i<30; i+=2) {
      {
         short sum=0;
         for (j=0; j<40; j++) {
             sum += a[j][i];
      }
         c[i] = sum;
      }
      {
         short sum=0;
         for (j=0; j<40; j++) {
             sum += a[j][i];
      }
         c[i+1] = sum;
      }
   }
}
```

The two inner loops can be jammed together. We shall assume that we have a `plus_eq2` operation which is a more powerful version of += that can handle two integers at a time. The result is:

```
void f_unroll_and_jam(int a[][40], short *restrict c) {
   int i, j;

   __builtin_aligned(a, 4);
   __builtin_aligned(c, 4);
   for (i=0; i<30; i+=2) {
```

```
        short sum0=0;
        short sum1=0;
        for (j=0; j<40; j++) {
            (sum0, sum1) .plus_eq2. (a[j][i], a[j][i+1]);
        }
        (c[i], c[i+1]) = (sum0, sum1);
    }
}
```

The above sequence of transformation, where an outer loop is unrolled and the copies corresponding to the inner loops are jammed together is called unroll and jam.

Assembly Code Unroll and Jam Example:

The assembly-annotated code for the above `f_unroll_and_jam` example is:

```
                M0 =  80 (X):
                LSETUP (._P1L2 , ._P1L6-2) LC1 = P2;
//  "f_unroll_and_jam.c" line 8 col 83
                P2 =   39;

.P1L2:
//-------------------------------------------------------------
//  Loop at "f_unroll_and_jam.c" line 6 col 4
//-------------------------------------------------------------
//  Loop was unrolled for unroll and jam 2 times
//-------------------------------------------------------------

                I0 = P0 :
                R0 = ROT R1 by 0  || NOP || R2 = [I0++M0];
                LSETUP (._P1L4 , ._P1L5-8) LC0 = P2;

.P1L4:
//-------------------------------------------------------------
//  Loop at "f_unroll_and_jam.c" line 8 col 8;
//-------------------------------------------------------------
//  This jammed loop executes 2 iterations of the original loop
//  in 1 cycle.
//  (1 iteration of the inner loop for each of the 2 unrolled
//  iterations of the outer loop)
//-------------------------------------------------------------
```

```
// This loop's resource usage is:
//    32-bit ALU/MAC used 1 out of 1 (100.0%)
//    Group 1 or 2 used 1 out of 2 ( 50.0%)
//--------------------------------------------------------------
//  Loop was jammed by unroll and jam 2 times
//--------------------------------------------------------------
//" f_unroll_and_jam.c" line 9 col 13
                R0 = R0 +|+ R2 || NOP || R2 = [I0++M0];
//" f_unroll_and_jam.c" line 8 col 8
                R0 = R0 + R2;
//      end loop ._P1L4;
//--------------------------------------------------------------
// End Loop L4
//--------------------------------------------------------------

.P1L5:
//--------------------------------------------------------------
//  Part of Loop 2, depth 1
//--------------------------------------------------------------
//" f_unroll_and_jam.c" line 9 col 13
                R0 = R0 +|+ R2;
//" f_unroll_and_jam.c" line 11 col 8
                [P1++] = R0;
                P0 +=  4;
//" f_unroll_and_jam.c" line 6 col 4
//      end loop ._P1L2;
//--------------------------------------------------------------
// End Loop L2
//--------------------------------------------------------------
```

Loop Flattening

Another transformation, slightly related to vectorization, is loop flattening. The loop flattening operation takes two nested loops that run `N1` and `N2` times, respectively, and transforms them into a single loop that runs `N1*N2` times.

For instance, the following function

```
void copy_v(int a[][100], int b[][100]) {
   int i,j;
   for (i=0; i< 30; ++i)
      for (j=0; j < 100; ++j)
          a[i][j] = b[i][j];
}
```

is transformed into

```
void copy_v(int a[][100], int b[][100]) {
   int i,j;
   int *p_a = &a[0][0];
   int *p_b = &b[0][0];
   for (i=0; i< 3000; ++i)
       p_a[i] = p_b[i];
}
```

This may further facilitate the vectorization process:

```
void copy_v(int a[][100], int b[][100]) {
   int i,j;
   int *p_a = &a[0][0];
   int *p_b = &b[0][0];
   for (i=0; i< 3000; i+=2)
       (p_a[i], p_a[i+1]) = (p_b[i], p_b[i+1]);
}
```

The assembly output for the loop flattening example is:

```
_copy_v:
//------------------------------------------------------------
........ procedure statistics ...........
//------------------------------------------------------------
-----
// Original Loop at "test_flatten_loop_dim.c" line 4 col 4 -- loop
   flattened into Loop at "test_flatten_loop_dim.c" line 5 col 7
//------------------------------------------------------------
..................... procedure code ..................

._P1L2:
```

```
//-------------------------------------------------------------
//   Loop at "test_flatten_loop_dim.c" line 5 col 7
//-------------------------------------------------------------
..................... loop annotations ...........
........................ loop body ...........
//-------------------------------------------------------------
//   End Loop L2
//-------------------------------------------------------------

._P1L3:
//-------------------------------------------------------------
//   Part of top level (no loop)
//-------------------------------------------------------------
// "test_flatten_loop_dim.c" line 7 col 1
                  RTS;
._copy_v.end:
```

Vectorization Annotations

For every loop that is vectorized, the following information is provided:

- The vectorization factor
- The number of peeled iterations
- The position of the peeled iterations (front or back of the loop)
- Information about whether peeled iterations are conditionally or unconditionally executed

For every loop pair that is subject to unroll and jam, it is necessary to:

- Annotate the unrolled outer loop with the number of times it was unrolled
- Annotate the inner loop with the number of times the loop was jammed

For every loop pair that is subject to loop flattening, it is necessary to account for the loop that is lost and show the remaining loop that it was merged with.

Example D:

Consider the test program:

```
void add(short *a, short *restrict b, short *restrict c, int dim) {
   int i, j;
   for (i = 0 ; i < dim; ++i)
       a[i] = b[i] + c[i];
}
```

for which the annotations produced are:

```
_add:
//------------------------------------------------------------
........ procedure statistics ...........
........ loop selection code  ...........
//------------------------------------------------------------
// Loop at "exampleE.c" line 4 col 4
//------------------------------------------------------------
// This loop executes 2 iterations of the original loop in 2 cycles.
//------------------------------------------------------------
................. loop resource annotation  .........
//------------------------------------------------------------
// Loop was vectorized by a factor of 2.
//------------------------------------------------------------
// Vectorization peeled 1 conditional iteration from the back
// of the loop because of an unknown trip count, possible
// not a multiple of 2.
//------------------------------------------------------------
....................... loop body .........
//------------------------------------------------------------
//   End Loop L8
//------------------------------------------------------------

//------------------------------------------------------------
..................... procedure code ...........
//------------------------------------------------------------
```

```
// Loop at "exampleE.c" line 4 col 4 (unvectorized version)
//-------------------------------------------------------------
// This loop executes 1 iteration of the original loop in 2 cycles.
//-------------------------------------------------------------
..................  loop resource annotation  .........
//-------------------------------------------------------------
.......................  loop body .........
//-------------------------------------------------------------
//   End Loop L11
```

The compiler has generated two versions of the loop—a vectorized version and a non-vectorized version. The vectorized version of the loop will be executed as long as all the pointers are sufficiently aligned. The compiler has peeled a single iteration from the end of the vectorized version of the loop, which will be executed if the pointers are all aligned, but `dim` is not a multiple of two.

3 C/C++ RUN-TIME LIBRARY

The C and C++ run-time libraries are collections of functions, macros, and class templates that are called from your source programs. The libraries provide a broad range of services including those that are basic to the languages such as memory allocation, character and string conversions, and math calculations. Using the library simplifies your software development by providing code for a variety of common needs.

This chapter contains:

- "C and C++ Run-Time Library Guide" on page 3-3
  provides introductory information about the ANSI/ISO standard C and C++ libraries. It also provides information about the ANSI standard header files and built-in functions that are included with this release of the `ccblkfn` compiler.
- "Documented Library Functions" on page 3-49
  tabulates the functions that are defined by ANSI standard header files.
- "C Run-Time Library Reference" on page 3-53
  provides reference information about the C run-time library functions included with this release of the `ccblkfn` compiler.

The `ccblkfn` compiler provides a broad collection of library functions, including those required by the ANSI standard and additional functions supplied by Analog Devices that are of value in signal processing applications. In addition to the standard C library, this release of the compiler software includes the abridged C++ library, a conforming subset of the standard C++ library. The abridged C++ library includes the embedded C++ and embedded standard template libraries.

This chapter describes the standard C/C++ library functions that are supported in the current release of the run-time libraries. Chapter 4, "DSP Run-Time Library", describes a number of signal processing, vector, matrix, and statistical functions that assist DSP code development.

For more information on the C standard library, see The Standard C Library by P.J. Plauger, Prentice Hall, 1992. For more information on the algorithms on which many of the C library's math functions are based, see Cody, W. J. and W. Waite, *Software Manual for the Elementary Functions*, Englewood Cliffs, New Jersey: Prentice Hall, 1980. For more information on the C++ library portion of the ANSI/ISO Standard for C++, see Plauger, P. J. (Preface), *The Draft Standard C++ Library*, Englewood Cliffs, New Jersey: Prentice Hall, 1994, (ISBN: 0131170031).

The C++ library reference information in HTML format is included on the software distribution CD-ROM. To access the reference files from VisualDSP++, use the **Help Topics** command (**Help** menu) and select the **Reference** book icon. From the **Online Manuals** topic, you can open any of the library files.

C and C++ Run-Time Library Guide

The C/C++ run-time libraries contain functions that can be called from your source. This section describes how to use the library and provides information on these topics:

- "Calling Library Functions" on page 3-3
- "Using the Compiler's Built-In Functions" on page 3-5
- "Linking Library Functions" on page 3-5
- "Library Function Re-Entrancy and Multi-Threaded Environments" on page 3-9
- "Working With Library Header Files" on page 3-14
- "Calling Library Functions from an ISR" on page 3-26
- "Abridged C++ Library Support" on page 3-27
- "File I/O Support" on page 3-34

For information on the C library's contents, see "C Run-Time Library Reference" on page 3-53. For information on the Abridged C++ library's contents, see "Abridged C++ Library Support" on page 3-27 and the VisualDSP++ online Help.

Calling Library Functions

To use a C/C++ library function, call the function by name and give the appropriate arguments. The names and arguments for each function appear on the function's reference page. These reference pages appear in "C Run-Time Library Reference" on page 3-53.

Like other functions, library functions should be declared. Declarations are supplied in header files, as described in "Working With Library Header Files" on page 3-14.

Function names are C/C++ function names. If you call a C or C++ run-time library function from an assembler program, you must use the assembly version of the function name.

- For C functions, this is an underscore at the beginning of the C function name. For example, the C function `main()` is referred to as `_main` from an assembler program.
- Functions in C++ modules are normally compiled with an encoded function name. Function names in C++ contain abbreviations for the parameters to the function and also the return type. As such, they can become very large. The compiler "mangles" these names to a shorter form. You can instruct the C++ compiler to use the single-underscore convention from C, as shown by the following example.

```
extern "C" {
    int myfunc(int);   // external name is _myfunc
}
```

Alternatively, compile C++ files to assembler, and see how the function has been declared in the assembly file.

It may not be possible to call inline functions as the compiler may have removed the definition of the function if all calls to the function are inlined. Global static variables cannot be referred to in assembly routines as their names are encrypted.

For more information on naming conventions, see "C/C++ and Assembly Interface" on page 1-269.

Use the archiver (`elfar`), described in the *VisualDSP++ 4.0 Linker and Utilities Manual*, to build library archive files of your own functions.

Using the Compiler's Built-In Functions

The C/C++ compiler's built-in functions are a set of functions that the compiler immediately recognizes and replaces with inline assembly code instead of a function call. Typically, inline assembly code is faster than a library routine, and it does not incur the calling overhead. For example, the absolute value function, `abs()`, is recognized by the compiler, which subsequently replaces a call to the C/C++ run-time library version with an inline version.

To use built-in functions, include the appropriate headers in your source, otherwise, your program build is going to fail at link time. If you want to use the C/C++ run-time library functions of the same name, compile using the `-no-builtin` compiler switch (on page 1-41).

Standard math functions, such as abs, min, max, are implemented using compiler built-in functions. They perform as documented in "C Run-Time Library Reference" on page 3-53 and "DSP Run-Time Library Reference" on page 4-43.

Linking Library Functions

The C/C++ run-time library is organized as a set of run-time libraries and startup files that are installed under the VisualDSP++ installation directory in the subdirectory `Blackfin\lib`. Table 3-1 contains a list of these library files together with a brief description of their functions.

Table 3-1. C and C++ Library Files

Blackfin\lib Directory	Description
`crt*.doj`	C run-time startup file which sets up system environment before calling `main()`
`crtn*.doj`	C++ cleanup file used for C++ constructors and destructors

Table 3-1. C and C++ Library Files (Cont'd)

Blackfin\lib Directory	Description
cplbtab*.doj	Default cache configuration table; memory protection and caching attributes for each Blackfin processor's memory map. See "Caching and Memory Protection" on page 1-208.
idle*.doj	Normal "termination" code that enters IDLE loop after "end" of the application
__initsbsz*.doj	Memory initializer support files
libc*.dlb	Primary ANSI C run-time library
libcpp*.dlb	Primary ANSI C++ run-time library
libcpprt*.dlb	C++ run-time support library
libdsp*.dlb	DSP run-time library
libetsi*.dlb	ETSI run-time support library
libio*.dlb	Host-based I/O facilities, as described in"stdio.h" on page 3-21
libevent*.dlb	Interrupt handler support library
libx*.dlb	C++ exception handling support library
lib*.dlb	C/C++ run-time library routines that save context, such as setjmp, longjmp, and exception handlers.
libprofile*.dlb	Profile support routines
librt*.dlb	C run-time support library; without File I/O
librt_fileio*.dlb	C run-time support library, with File I/O
libsftflt*.dlb	Floating-point emulation routines
libsmall*.dlb	Supervisor mode support routines
prfflg0*.doj prfflg1*.doj prfflg2*.doj	Profiling initialization routines as selected by -p, -p1, and -p2 compiler options (see "-p[1|2]" on page 1-47)

In general, several versions of each C/C++ run-time library component is supplied in binary form; for example, variants are available for different Blackfin architectures while other variants have been built for running in a multi-threaded environment. Each version of a library or startup file is distinguished by a different combination of filename suffices.

Table 3-2 lists the filename suffices that may be used.

Table 3-2. Filename Suffices

Filename Suffix	Description
`531`	Compiled only for ADSP-BF531 processor
`532`	Compiled for execution on ADSP-BF531, ADSP-BF532, ADSP-BF533, or ADSP-BF561 processors
`533`	Compiled only for ADSP-BF533 processor
`535`	Compiled for execution on a ADSP-BF535 or AD6532 processor
`561`	Compiled only for ADSP-BF561 processor
`a`	Compiled only for ADSP-BF561 Core A processor
`b`	Compiled only for ADSP-BF561 Core B processor
`mt`	Built for multi-thread environments
`x`	Libraries that are compiled with C++ exception handling enabled
`y`	Compiled with the `-workaround all` switch (on page 1-62) to avoid all known hardware anomalies
`m3res`	Libraries that are compiled with the `-reserve M3` switch and are usually used for running under an emulator (see Emulator documentation)
`m3free`	Libraries that may use the M3 register and may not therefore be suitable for running on an emulator

For example, the C run-time library `libc535mty.dlb` has been compiled with the `-workaround all` switch (see on page 1-62) for execution on either an ADSP-BF535 or AD6532 processor, and has been built for the VDK multi-threaded environments.

The C/C++ run-time library provides further variants of the startup files (`crt*.doj`) that have been built from a single source file (see "Startup Code Overview" on page 1-196). Table 3-3 shows the filename suffices that are used to differentiate between different versions of this binary file.

Table 3-3. crt Filename Suffices

crt Filename Suffix	Description
c	Startup file used for C++ applications
f	Startup file that enables file I/O support via `stdio.h`
p	Startup file used by applications that have been compiled with profiling instrumentation
s	Startup file used by applications that run in supervisor mode

For example, the file `crtcf535.doj` is the startup file that enables file I/O support and initializes a C++ application that has been compiled to execute in user mode on either an ADSP-BF535 or AD6532 processor.

When an application calls a C or C++ library function, the call creates a reference that the linker resolves. One way to direct the linker to the location of the appropriate run-time library is to use the default Linker Description File (`<your_target >.ldf`). If a customized `.LDF` file is used to link the application, then the appropriate library/libraries and startup files should be added to the `.LDF` file used by the project.

Instead of modifying a customized `.LDF` file, the compiler's `-l` switch can be used to specify which libraries should be searched by the linker. For example, the switches `-lc532 -lcpp532 -lcpprt532` add the C library `libc532.dlb` as well as the C++ libraries `libcpp532.dlb` and `libcpprt532.dlb` to the list of libraries that the linker examines. For more information on the `.LDF` file, see the *VisualDSP++ 4.0 Linker and Utilities Manual*.

Library Function Re-Entrancy and Multi-Threaded Environments

C/C++ run-time libraries are not re-entrant. For example, it is not possible to put the library code into L2 memory on the ADSP-BF561 processor and have either core (Core A or Core B) to call the libraries without using some user-defined semaphore.

It is sometimes desirable for there to be several active instances of a given library function at once. This can occur because:

- An interrupt or other external event invokes a function, while the application is also executing that function.
- The application uses a multi-threaded architecture, such as VDK, and more than one thread executes the function concurrently.
- The application is built for a multi-core processor, such as the ADSP-BF561 processor, and more than one core is executing the function concurrently.

When there may be multiple concurrent threads active at once, it is important to ensure that the library functions used are able to support this activity; if a function uses private data storage and both active instances of the function modify the same storage area without due care, undefined behavior may occur.

The majority of the C/C++ run-time library is safe in this regard, in that the functions do not have private storage, operating instead on parameters passed by the caller.

A small subset of Standard functions make use of private storage, and functions like the `stdio` support operate on shared resources (like `FILE` pointers) that need to be safely updated. To support these needs, multi-threaded builds of the run-time libraries are included in VisualDSP++.

The multi-threaded versions of the run-time libraries make use of local-storage routines, for thread-local and core-local private copies of data (see "adi_obtain_mc_slot" on page 3-60 for more information), and have recursive locking mechanisms included so that shared resources, such as `stdio FILE` buffers, are only updated by a single function instance at any given time.

There is a minor difference between "multi-threaded" and "multi-core":

- The multi-threaded libraries have "`mt`" in their filename, and are built for an arbitrary number of concurrent threads, as is the case for VDK applications. Multi-threaded libraries are used both for VDK builds and for non-VDK builds on dual-core processors, when the `-multicore` compiler switch (on page 1-40) is specified.
- The multi-core libraries have "`mc`" in their filename, and are built for dual-threaded applications, with a single thread running on each of two cores in a dual-core processor. The multi-core libraries are used for non-VDK builds on dual-core processors, when the `-multicore` compiler switch is specified.

The following Standard library elements make use of thread-local or core-local private storage:

- the `atexit` function
- the `rand` function
- the `strtok` function
- the `asctime` function
- the `errno` global variable

The `atexit` function requires a core-local slot, but not thread-local slot. This is because VDK applications do not use `exit` to terminate each thread, and effectively run "forever"; using `exit` terminates the whole VDK application. In contrast, in a multi-core application, `exit` could terminate the application in one core while the other continues.

You must specify the -multicore compiler switch when building multi-core applications. The switch has the following effects:

- At compile-time, it defines the `__ADI_MULTICORE` macro to ensure that core-local storage operations are available.
- At link-time, it ensures that the application is linked against the multi-threaded and multi-core builds of the library.
- It repositions the default heap to be in shared memory. Allocations by either core will be served by the same heap. The heap allocation and release routines use locking to ensure that only one core is updating the heap resource records at once.

While thread-safe variants of the C/C++ run-time libraries exist, many functions are not interrupt-safe as they access global data structures. It is therefore recommended that ISRs do not make library function calls, as unexpected behavior may result if the interrupt occurs during a call to such a function. An alternative approach is to disable interrupts before the application makes run-time library calls. This may be disadvantageous for time-critical applications as interrupts may be disabled for a long period of time. The DSP run-time library functions do not modify global data structures and are therefore interrupt-safe.

Support Functions for Private Data

The run-time library provides support functions for creating thread-local and core-local private data storage.

For thread-local private storage, refer to the *VisualDSP++ Kernel (VDK) User's Guide*.

For core-local private storage, see "adi_obtain_mc_slot" on page 3-60.

Support Functions for Locking

The run-time library provides support functions for ensuring safe access to shared resources, in the form of locking routines. See "adi_acquire_lock" on page 3-57 for more information.

Other Support Functions for Multi-Core Applications

The run-time library includes the `adi_core_id` function, which can be used by shared code to determine which core is executing it. See "adi_core_id" on page 3-59 for more details.

Library Placement

A multi-threaded or multi-core application has some storage that must be shared across threads and cores (such as locks, that must be globally-accessible), and some storage that must be private (such as the C++ exception look-up tables in a multi-core application). The following table lists the requirements for each of the libraries, regarding section placement.

Object/library multicore restrictions are:

Object/Library	Restriction
`__init*.doj`, `bootup*.doj`, `cplbtab*.doj`, `crt*.doj`	The startup and configuration objects are all core-specific, and must not be placed, in a shared memory section
`libc*.dlb`	Can be placed in a shared memory section
`libcpp*.dlb`, `libcpprt*.dlb`	Cannot be placed in a shared memory section
`libdsp*.dlb`	Can be placed in a shared memory section

Object/Library	Restriction
libetsi*.dlb	Can be placed in a shared memory section, as provided. If re-built in debug mode, so that Overflow and Carry "flags" are maintained, these global variables will not be locked during updates.
libevent*.dlb	Single core memory placement recommended
libm3*.dlb	Can be placed in a shared memory section
libmc*.dlb	Can be placed in a shared memory section
mc_data*.doj	Must be placed in a shared memory section
libprofile*.dlb	Not supported in a multi-core environment
librt*.dlb	Can be placed in a shared memory section
librt_fileio.dlb	Can be placed in a shared memory section
libsftflt*.dlb	Can be placed in a shared memory section
libsmall*.dlb	Can be placed in a shared memory section
libx*.dlb	Cannot be placed in a shared memory section
prfflg*.doj	Not supported in a dual-core environment

Section Placement

Libraries are mapped into shared or private areas through the use of sections within the .LDF file. The following table shows which LDF output sections must be shared, and which must be private.

LDF Section	Must Be Shared	Must Be Core-Specific
primio_atomic_lock	Yes	No
mc_data	Yes	No
heap	No	No
L1_code	No	Yes
cplb	No	Yes
cplb_data	No	Yes
bsz	No	Yes

LDF Section	Must Be Shared	Must Be Core-Specific
bsz_init	No	Yes
.edt	No	Yes
.cht	No	Yes
constdata	No	Yes
ctor	No	Yes
ctorl	No	Yes
.gdt	No	Yes
.gdtl	No	Yes
.frt	No	Yes
.frtl	No	Yes
stack	No	Yes

Note that the sharing or privacy of the "heap" section is a matter for the application designer. A multi-core application defaults to using a shared heap with appropriate locking during allocation and release, but a private per-core heap may better suit the application.

Any sections not listed above may be shared or private, according to the discretion of the application developer. You must ensure that any shared sections use appropriate locking mechanisms to avoid corruption by simultaneous accesses.

Working With Library Header Files

When you use a library function in your program, you should also include the function's header file with the `#include` preprocessor command. The header file for each function is identified in the *Synopsis* section of the function's reference page. Header files contain function prototypes. The compiler uses these prototypes to check that each function is called with the correct arguments.

A list of the header files that are supplied with this release of the Blackfin compiler appears in Table 3-4. Use a C standard text to augment the information supplied in this chapter.

Table 3-4. Standard C Run-Time Library Header Files

Header	Purpose	Standard
assert.h	Diagnostics	ANSI
ctype.h	Character Handling	ANSI
device.h	Macros and data structures for alternative device drivers	Analog extension
device_int.h	Enumerations and prototypes for alternative device drivers	Analog extension
errno.h	Error Handling	ANSI
float.h	Floating Point	ANSI
iso646.h	Boolean Operators	ANSI
limits.h	Limits	ANSI
locale.h	Localization	ANSI
math.h	Mathematics	ANSI
setjmp.h	Non-Local Jumps	ANSI
signal.h	Signal Handling	ANSI
stdarg.h	Variable Arguments	ANSI
stddef.h	Standard Definitions	ANSI
stdio.h	Input/Output	ANSI
stdlib.h	Standard Library	ANSI
string.h	String Handling	ANSI
time.h	Date and Time	ANSI

The following sections provide descriptions of the header files contained in the C library. The header files are listed in alphabetical order.

assert.h

The `assert.h` file contains the `assert` macro.

ctype.h

The `ctype.h` file contains functions for character handling, such as `isalpha`, `tolower`, etc.

device.h

The `device.h` header file provides macros and defines data structures that an alternative device driver would require to provide file input and output services for `stdio` library functions. Normally, the `stdio` functions use a default driver to access an underlying device, but alternative device drivers may be registered that may then be used transparently by these functions. This mechanism is described in "Extending I/O Support To New Devices" on page 3-35.

device_int.h

The `device_int.h` header file contains function prototypes and provides enumerations for alterative device drivers. An alternative device driver is normally provided by an application and may be used by the `stdio` library functions to access an underlying device; an alternative device driver may coexist with, or may replace, the default driver that is supported by the VisualDSP++ simulator and EZ-KIT Lite evaluation systems. Refer to "Extending I/O Support To New Devices" on page 3-35 for more information.

errno.h

The `errno.h` file provides access to `errno`. This facility is not, in general, supported by the rest of the library.

float.h

The `float.h` header file defines the properties of the floating-point data types that are implemented by the compiler—that is, `float`, `double`, and `long double`. These properties are defined as macros and include the following for each supported data type:

- the maximum and minimum value (for example, `FLT_MAX` and `FLT_MIN`)
- the maximum and minimum power of ten (for example, `FLT_MAX_10_EXP` and `FLT_MIN_10_EXP`)
- the precision available expressed in terms of decimal digits (for example, `FLT_DIG`)
- a constant that represents the smallest value that may added to 1.0 and still result in a change of value. (for example, `FLT_EPSILON`)

Note that the set of macros that define the properties of the `double` data type will have the same values as the corresponding set of macros for the `float` type when `doubles` are defined to be 32 bits wide, and they will have the same value as the macros for the `long double` data type when `doubles` are defined to be 64 bits wide (see "-double-size-{32 | 64}" on page 1-29).

iso646.h

The `iso646.h` header file defines symbolic names for certain C (Boolean) operators; the symbolic names and their associated value are shown in Table 3-5.

The symbolic names have the same name as the C++ keywords that are accepted by the compiler when the `-alttok` switch is specified (see on page 1-26).

Table 3-5. Symbolic Names Defined in iso646.h

Symbolic Name	Equivalent
and	&&
and_eq	&=
bitand	&
bitor	|
compl	~
not	!
not_eq	!=
or	||
or_eq	|=
xor	^
xor_eq	^=

limits.h

The `limits.h` file contains definitions of maximum and minimum values for each C data type other than a floating-point type.

locale.h

The `locale.h` file contains definitions for expressing numeric, monetary, time, and other data.

math.h

The `math.h` header file includes trigonometric, power, logarithmic, exponential, and other miscellaneous functions. The library contains the functions specified by the C standard along with implementations for the data types `float` and `long double`. Some functions also support a 16-bit fractional data type.

For every function that is defined to return a `double`, the `math.h` header file also defines two corresponding functions that return a `float` and a `long double`, respectively. The names of the `float` functions are the same as the equivalent `double` function with `f` appended to its name. Similarly, the names of the `long double` functions are the same as the `double` function with `d` appended to its name. For example, the header file contains the following prototypes for the sine function:

```
float sinf (float x);
double sin (double x);
long double sind (long double x);
```

The `-double-size-{32|64}` compiler switch (on page 1-29) controls the size of the `double` data type. The default behavior is for the compiler to compile the `double` type as a 32-bit floating point data type, and the header file will arrange that all references to a `double` function is directed to the equivalent float function (with the suffix `f`). Conversely, when the `double` type is defined as a 64-bit floating point data type, all references to the `double` functions of this header file will be directed to the `long double` version of the function (with the suffix `d`). This allows the un-suffixed function names to be used with arguments of type `double`, regardless of whether `doubles` are 32 or 64 bits long.

The `math.h` file also defines the macro `HUGE_VAL` which evaluates to the maximum positive value that the type `double` can support.

Some of the functions in this header file exist as both integer and floating point. The floating-point functions typically have an `f` prefix. Make sure you are using the correct one.

The C language provides for implicit type conversion, so the following sequence produces surprising results with no warnings.

```
float x,y;
y = abs(x);
```

The value in `x` is truncated to an integer prior to calculating the absolute value, then reconverted to floating point for the assignment to `y`.

setjmp.h

The `setjmp.h` file contains `setjmp` and `longjmp` for non-local jumps.

signal.h

The `signal.h` file provides function prototypes for the standard ANSI `signal.h` routines. It also includes ANSI standard signal handling functions of the C library.

The signal handling functions process conditions (hardware signals) that can occur during program execution. They determine the way C programs respond to these signals. The functions are designed to process such signals as external interrupts and timer interrupts.

stdarg.h

The `stdarg.h` header file contains definitions needed for functions that accept a variable number of arguments. Programs that call such functions must include a prototype for the functions referenced.

stddef.h

The `stddef.h` file contains a few common definitions useful for portable programs, such as `size_t`.

stdio.h

The `stdio.h` header file defines a set of functions, macros, and data types for performing input and output. The library functions defined by this header file are thread-safe but they are not in general interrupt-safe, and therefore they should not be called either directly or indirectly from an interrupt service routine.

The compiler uses the definitions within the header file to select an appropriate set of functions that correspond to the currently selected size of type `double` (either 32 bits or 64 bits). Any source file that uses the facilities of `stdio.h` should therefore include the header file, especially if it is compiled with the `-double-size-64` switch (on page 1-29). Failure to include the header file may result in a linker failure as the compiler must see a correct function prototype in order to generate the correct calling sequence.

This release provides two alternative run-time libraries that implement the functionality of the header file. If an application is built with the `-full-io` switch (on page 1-33), then it is linked with a third-party I/O library that provides a comprehensive implementation of the ANSI C Standard I/O functionality, but at the cost of some performance. This library is fully compatible with previous releases of VisualDSP++ (version 3.5 and earlier). No source files are provided for this proprietary library.

However, the normal behavior of the compiler is to link an application against an I/O library provided by Analog Devices—this library does not support all the facilities of the third party library, but it is both faster and smaller. The source files for this library are available under the VisualDSP++ installation in the subdirectory `Blackfin\lib\src\libio`.

The functional differences between the library based on third-party software (and accessed via the `-full-io` switch) and the default I/O run-time library provided by Analog Devices are given below:

- The third party I/O library supports the input and output of wide characters (that is, data of type `wchar_t`), and multi-byte characters. No similar support is available under the Analog Devices I/O library.
- The `fread()` and `fwrite()` functions are commonly used to transmit data between an application and a binary stream. The Analog Devices I/O library does not normally use a buffer to read or write data using these functions; thus, the data is usually transmitted directly between a program and an external device. If an application relies on these functions to read and write data via an I/O buffer, then it should be linked against the third-party library (using the `-full-io` switch).
- The third-party I/O library supports additional size qualifiers, which are defined in the C99 (ISO/IEC 9899:1999) standard; these additional qualifiers may all be used with the `d`, `i`, `o`, `u`, `x`, or `X` conversion specifiers to describe the type of the corresponding argument, such as:

`hh`	`signed char` ***or*** `unsigned char`
`j`	`intmax_t` ***or*** `uintmax_t`
`t`	`ptrdiff_t`
`z`	`size_t`

- The third-party I/O library accesses the current locale to determine the symbol to be used as the decimal point character.
- The alternative libraries have different conventions for printing IEEE floating-point values that are either NaN's (Not-A-Number) or Infinity. The third party I/O library also accepts `nan` and `inf` (in any case) as input for the `e`, `f`, and `g` conversion specifiers.
- The form of the output generated for the a conversion specifier by the alternative libraries differ (both forms of output do however conform to the requirements of ISO/IEC 9899:1999).

- The conversion specifier `F` is accepted by the third-party I/O library; the specifier behaves the same as `f`.
- The third party I/O library also supports the full functionality of the `[` conversion specifier, while the Analog Devices I/O library only provides the minimum facility level required by the ANSI standard.

The implementation of both I/O libraries is based on a simple interface with a device driver that provides a set of low-level primitives for `open`, `close`, `read`, `write`, and `seek` operations. By default, these operations are provided by the VisualDSP++ simulator and EZ-KIT Lite systems and this mechanism is outlined in "Default Device Driver Interface" on page 3-43. However, alternative device drivers may be registered (see "Extending I/O Support To New Devices" on page 3-35) that can then be used transparently through the `stdio.h` functions.

The following restrictions apply to either library in this software release:

- Functions `tmpfile()` and `tmpnam()` are not available.
- Functions `rename()` and `remove()` are only supported under the default device driver supplied by the VisualDSP++ simulator and EZ-KIT Lite system, and they only operate on the host file system.
- Positioning within a file that has been opened as a text stream is only supported if the lines within the file are terminated by the character sequence `\r\n`.
- Support for formatted reading and writing of data of type `long double` is only supported if an application is built with the `-double-size-64` switch.

stdlib.h

The `stdlib.h` header file offers general utilities specified by the C standard. These include some integer math functions, such as `abs`, `div`, and `rand`; general string-to-numeric conversions; memory allocation functions, such as `malloc` and `free`; and termination functions, such as `exit`. This library also contains miscellaneous functions such as `bsearch` and `qsort`.

string.h

The `string.h` file contains string handling functions, including `strcpy` and `memcpy`.

time.h

The `time.h` header file provides functions, data types, and a macro for expressing and manipulating date and time information. The header file defines two fundamental data types, one of which is `clock_t` and is associated with the number of implementation-dependent processor "ticks"' used since an arbitrary starting point; and the other which is `time_t`.

The `time_t` data type is used for values that represent the number of seconds that have elapsed since a known epoch; values of this form are known as a calendar time. In this implementation, the epoch starts on 1st January, 1970, and calendar times before this date are represented as negative values.

A calendar time may also be represented in a more versatile way as a broken-down time. A broken-down time is a structured variable of the following form:

```
struct tm { int tm_sec;  /* seconds after the minute [0,61] */
  int tm_min;           /* minutes after the hour [0,59]   */
  int tm_hour;          /* hours after midnight [0,23]     */
  int tm_mday;          /* day of the month [1,31]         */
  int tm_mon;           /* months since January [0,11]     */
  int tm_year;          /* years since 1900                */
```

```
    int tm_wday;          /* days since Sunday [0, 6]        */
    int tm_yday;          /* days since January 1st [0,365]  */
    int tm_isdst;         /* Daylight Saving flag            */
};
```

This implementation does not support either the Daylight Saving flag in the structure `struct tm`; nor does it support the concept of time zones. All calendar times are therefore assumed to relate to Greenwich Mean Time (Coordinated Universal Time or UTC).

The header file sets the macro `CLOCKS_PER_SEC` to the number of processor cycles per second and this macro can therefore be used to convert data of type `clock_t` into seconds, normally by using floating-point arithmetic to divide it into the result returned by the `clock` function.

In general, the processor speed is a property of a particular chip and it is therefore recommended that the value to which this macro is set is verified independently before it is used by an application.

On Blackfin processors, the macro is set by one of the following (in descending order of precedence):

- Via the `-DCLOCKS_PER_SEC=<definition>` compile-time switch. Because the `time_t` type is based on the `long long int` data type, it is recommended that the value of the symbolic name `CLOCKS_PER_SEC` is defined to be of type `long long int` by qualifying the value with the `LL` (or `ll`) suffix (that is, `-DCLOCKS_PER_SEC=6000000LL`).
- Via the System Services Library
- Via the **Processor speed** box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **Processor** (**1**) category
- From the header file `cycles.h`

Calling Library Functions from an ISR

Not all C run-time library functions are interrupt-safe (and can therefore be called from an Interrupt Service Routine). For a run-time function to be classified as *interrupt-safe*, it must:

- not update any global data, such as `errno`, and
- not write to (or maintain) any private static data

It is recommended therefore that none of the functions defined in the header file `math.h`, nor the string conversion functions defined in `stdlib.h`, be called from an ISR as these functions are commonly defined to update the global variable `errno`. Similarly, the functions defined in the `stdio.h` header file maintain static tables for currently opened streams and should not be called from an ISR.

Several other library functions are not interrupt-safe because they make use of private static data. These functions are:

```
asctime
gmtime
localtime
rand
srand
strtok
```

While not all C run-time library functions are interrupt-safe, versions of the functions are available that are *thread-safe* and may be used in in either a VDK multi-threaded environment or by dual-core Blackfin applications (see "Reentrancy" on page A-18). These library functions can be found in the run-time libraries that have the suffix `_mt` in their filename.

Abridged C++ Library Support

When in C++ mode, the compiler can call a large number of functions from the Abridged Library, a conforming subset of the C++ library.

The Abridged Library has two major components: embedded C++ library (EC++) and embedded standard template library (ESTL). The embedded C++ library is a conforming implementation of the embedded C++ library as specified by the Embedded C++ Technical Committee. For more information on the Abridged library, see online Help.

This section lists and briefly describes the following components of the Abridged library:

- "Embedded C++ Library Header Files" on page 3-27
- "C++ Header Files for C Library Facilities" on page 3-30
- "Embedded Standard Template Library Header Files" on page 3-31
- "Using the Thread-Safe C/C++ Run-Time Libraries with VDK" on page 3-34

Embedded C++ Library Header Files

The following section provides a brief description of the header files in the embedded C++ library.

complex

The `complex` header file defines a template class complex and a set of associated arithmetic operators. Predefined types include `complex_float` and `complex_long_double`.

This implementation does not support the full set of complex operations as specified by the C++ standard. In particular, it does not support either the transcendental functions or the I/O operators << and >>. The `complex` header and the C library header file `complex.h` refer to two different and incompatible implementations of the `complex` data type.

exception

The `exception` header file defines the `exception` and `bad_exception` classes and several functions for exception handling.

fract

The `fract` header file defines the `fract` data type, which supports fractional arithmetic, assignment, and type-conversion operations using a 32-bit data type. The header file is fully described under "Fractional Value Built-In Functions in C++" on page 1-128.

fstream

The `fstream` header file defines the `filebuf`, `ifstream`, and `ofstream` classes for external file manipulations.

iomanip

The `iomanip` header file declares several `iostream` manipulators. Each manipulator accepts a single argument.

ios

The `ios` header file defines several classes and functions for basic `iostream` manipulations. Note that most of the `iostream` header files include `ios`.

iosfwd

The `iosfwd` header file declares forward references to various `iostream` template classes defined in other standard headers.

iostream

The `iostream` header file declares most of the `iostream` objects used for the standard stream manipulations.

istream

The `istream` header file defines the `istream` class for iostream extractions. Note that most of the `iostream` header files include `istream`.

new

The `new` header file declares several classes and functions for memory allocations and deallocations.

ostream

The `ostream` header file defines the `ostream` class for `iostream` insertions.

shortfract

The `shortfract` header file defines the `shortfract` data type, which supports fractional arithmetic, assignment, and type-conversion operations using a 16-bit base type. The header file is fully described under "Fractional Value Built-In Functions in C++" on page 1-128.

sstream

The `sstream` header file defines the `stringbuf`, `istringstream`, and `ostringstream` classes for various `string` object manipulations.

stdexcept

The `stdexcept` header file defines a variety of classes for exception reporting.

streambuf

The `streambuf` header file defines the `streambuf` classes for basic operations of the `iostream` classes. Note that most of the iostream header files include `streambuf`.

string

The `string` header file defines the `string` template and various supporting classes and functions for `string` manipulations. Objects of the `string` type should not be confused with the null-terminated C strings.

strstream

The `strstream` header file defines the `strstreambuf`, `istrstream`, and `ostream` classes for iostream manipulations on allocated, extended, and freed character sequences.

C++ Header Files for C Library Facilities

For each C standard library header there is a corresponding standard C++ header. If the name of a C standard library header file is `foo.h`, then the name of the equivalent C++ header file will be `cfoo`. For example, the C++ header file `cstdio` provides the same facilities as the C header file `stdio.h`.

Table 3-6 lists the C++ header files that provide access to the C library facilities.

Normally, the C standard headers files may be used to define names in the C++ global namespace while the equivalent C++ header files define names in the standard namespace. However, the standard namespace is not supported in this release of the compiler. Therefore, the effect of including one of the C++ header files is the same as including the equivalent C standard library header file. When porting a source written to use namespace `std`, it may be necessary to exclude using statements and "`std::`" references. A useful way of removing `std::` references is to define `std` as a macro to nothing (for example, by using the switch `-Dstd=""`).

Table 3-6. C++ Header Files for C Library Facilities

Header	Description
cassert	Enforces assertions during function executions
cctype	Classifies characters
cerrno	Tests error codes reported by library functions
cfloat	Tests floating-point type properties
climits	Tests integer type properties
clocale	Adapts to different cultural conventions
cmath	Provides common mathematical operations
csetjmp	Executes non-local goto statements
csignal	Controls various exceptional conditions
cstdarg	Accesses a various number of arguments
cstddef	Defines several useful data types and macros
cstdio	Performs input and output
cstdlib	Performs a variety of operations
cstring	Manipulates several kinds of strings

Embedded Standard Template Library Header Files

Templates and the associated header files are not part of the embedded C++ standard library, but are supported by the compiler in C++ mode. The embedded standard template library headers are:

algorithm

The algorithm header defines numerous common operations on sequences.

deque

The deque header defines a deque template container.

functional

The `functional` header defines numerous function objects.

hash_map

The `hash_map` header defines two hashed map template containers.

hash_set

The `hash_set` header defines two hashed set template containers.

iterator

The `iterator` header defines common iterators and operations on iterators.

list

The `list` header defines a list template container.

map

The `map` header defines two map template containers.

memory

The `memory` header defines facilities for managing memory.

numeric

The `numeric` header defines several numeric operations on sequences.

queue

The `queue` header defines two queue template container adapters.

set

The `set` header defines two set template containers.

stack

The `stack` header defines a stack template container adapter.

utility

The `utility` header defines an assortment of utility templates.

vector

The `vector` header defines a vector template container.

The Embedded C++ library also includes several headers for compatibility with traditional C++ libraries, such as:

fstream.h

The `fstream.h` header defines several `iostreams` template classes that manipulate external files.

iomanip.h

The `iomanip.h` header defines several `iostreams` manipulators that take a single argument.

iostream.h

The `iostream.h` header declares the `iostreams` objects that manipulate the standard streams.

new.h

The `new.h` header declares several functions that allocate and free storage.

Using the Thread-Safe C/C++ Run-Time Libraries with VDK

When developing for VDK, the thread-safe variants of the run-time libraries are linked with user applications. These libraries may add an overhead to the VDK resources required by some applications.

The run-time libraries make use of VDK semaphores and unscheduled regions to ensure thread safety. The C and C++ libraries use a constant number of semaphores that are privately allocated by VDK, and these are in addition to any semaphores that a user may specify.

However the C++ `streambuf` class requires a semaphore for each instance of the class that is created. `Streambuf` objects are created as part of the C++ Standard I/O mechanism, and a `streambuf` may be created for each I/O stream that is created. VDK will privately allocate semaphores for the standard I/O streams `cout`, `cerr` and `cin`, and an additional semaphore is allocated for I/O accesses within the standard C library.

Users should be aware that they would need to increase the maximum semaphore count for their project if they create any additional `streambuf` objects.

File I/O Support

The VisualDSP++ environment provides access to files on a host system by using stdio functions. File I/O support is provided through a set of low-level primitives that implement the open, close, read, write, and seek operations. The functions defined in the `stdio.h` header file make use of these primitives to provide conventional C input and output facilities. The source files for the I/O primitives are available under the VisualDSP++ installation in the subdirectory `Blackfin\lib\src\libio`.

This section describes:

- "Extending I/O Support To New Devices" on page 3-35
- "Default Device Driver Interface" on page 3-43

Refer to "stdio.h" on page 3-21 for information about the conventional C input and output facilities that are provided by the compiler.

Extending I/O Support To New Devices

The I/O primitives are implemented using an extensible device driver mechanism. The default start-up code includes a device driver that can perform I/O through the VisualDSP++ simulator and EZ-KIT Lite evaluation systems. Other device drivers may be registered and then used through the normal `stdio` functions.

DevEntry Structure

A device driver is a set of primitive functions grouped together into a `DevEntry` structure. This structure is defined in `device.h`.

```
struct DevEntry {
   int   DeviceID;
   void  *data;

   int   (*init)(struct DevEntry *entry);
   int   (*open)(const char *name, int mode);
   int   (*close)(int fd);
   int   (*write)(int fd, unsigned char *buf, int size);
   int   (*read)(int fd, unsigned char *buf, int size);
   long  (*seek)(long fd, int offset, int whence);
   int stdinfd;
   int stdoutfd;
   int stderrfd;
}

typedef struct DevEntry DevEntry;
typedef struct DevEntry *DevEntry_t;
```

The fields within the `DevEntry` structure have the following meanings.

DeviceID:
The `DeviceID` field is a unique identifier for the device, known to the user. Device IDs are used globally across an application.

data:
The `data` field is a pointer for any private data the device may need; it is not used by the run-time libraries.

init:
The `init` field is a pointer to an initialization function. The run-time library calls this function when the device is first registered, passing in the address of this structure (and thus giving the `init` function access to `DeviceID` and the field `data`). If the `init` function encounters an error, it must return -1. Otherwise, it must return a positive value to indicate success.

open:
The `open` field is a pointer to a function performs the "*`open file`*" operation upon the device; the run-time library will call this function in response to requests such as `fopen()`, when the device is the currently-selected default device. The `name` parameter is the path name to the file to be opened, and the `mode` parameter is a bitmask indicating how the file is to be opened. Bits 0-1 indicate reading and/or writing.

```
0x0000   Open file for reading
0x0001   Open file for writing
0x0002   Open file for reading and writing
0x0003   Invalid
```

Additional bits may be ORed into `mode` to alter the file's behavior, such as:

```
0x0008   Open the file for appending
0x0100   Create the file, if it does not already exist.
0x0200   Truncate the file to zero length, if it
         already exists
```

```
0x4000   Open the file as a text stream (converting the
         character sequence \r\n to \n on reading,
         and \n to \r\n on writing).
0x8000   Open the file as a binary stream (raw mode).
```

The `open` function must return a positive "*`file descriptor`*" if it succeeds in opening the file; this file descriptor is used to identify the file to the device in subsequent operations. The file descriptor must be unique for all files currently open for the device, but need not be distinct from file descriptors returned by other devices—the run-time library identifies the file by the combination of device and file descriptor.

If the `open` function fails, it must return -1 to indicate failure.

close:
The `close` field is a pointer to a function that performs the "*`close file`*" operation on the device. The run-time library calls the `close` function in response to requests such as `fclose()` on a stream that was opened on the device. The `fd` parameter is a file descriptor previously returned by a call to the `open` function. The `close` function must return a positive value for success, and a negative value for failure.

write:
The `write` field is a pointer to a function that performs the "*`write to file`*" operation on the device. The run-time library calls the `write` function in response to requests, such as `fwrite()`, `fprintf()` and so on, that act on streams that were opened on the device. The `write` function takes three parameters:

- `fd` – this is a file descriptor that identifies the file to be written to; it will be a value that was returned from a previous call to the open function.
- `buf` – a pointer to the data to be written to the file
- `size` – the number of bytes to be written to the file

The `write` function must return one of the following values:

- A positive value from 1 to `size` inclusive, indicating how many bytes from `buf` were successfully written to the file
- Zero, indicating that the file has been closed, for some reason (for example, network connection dropped)
- A negative value, indicating an error

read:
The `read` field is a pointer to a function that performs the "*`read from file`*" operation on the device. The run-time library calls the `read` function in response to requests, such as `fread()`, `fscanf()` and so on, that act on streams that were opened on the device. The `read` function's parameters are:

- `fd` – this is the file descriptor for the file to be read
- `buf` – this is a pointer to the buffer where the retrieved data must be stored
- `size` – this is the number of (8-bit) bytes to read from the file. This must not exceed the space available in the buffer pointed to by `buf`.

The `read` function must return one of the following values:

- A positive value from 1 to `size` inclusive, indicating how many bytes were read from the file into `buf`
- Zero, indicating end-of-file
- A negative value, indicating an error

The run-time library expects the `read` function to return `0xa (10)` as the newline character.

seek:
The `seek` field is a pointer to a function that performs dynamic access on the file. The run-time library calls the `seek` function in response to requests such as `rewind()`, `fseek()`, and so on, that act on streams that were opened on the device.

The `seek` function takes the following parameters:

- `fd` – this is the file descriptor for the file which will have its read/write position altered
- `offset` – this is a value that is used to determine the new read/write pointer position within the file; it is in (8-bit) bytes
- `whence` – this is a value that indicates how the `offset` parameter is interpreted:
    - 0: `offset` is an absolute value, giving the new read/write position in the file
    - 1: `offset` is a value relative to the current position within the file
    - 2: `offset` is a value relative to the end of the file

The `seek` function returns a positive value that is the new (absolute) position of the read/write pointer within the file, unless an error is encountered, in which case the `seek` function must return a negative value.

If a device does not support the functionality required by one of these functions (such as read-only devices, or stream devices that do not support seeking), the `DevEntry` structure must still have a pointer to a valid function; the function must arrange to return an error for attempted operations.

stdinfd:
The `stdinfd` field is set to the device file descriptor for `stdin` if the device is expecting to claim the `stdin` stream, or to the enumeration value `dev_not_claimed` otherwise.

stdoutfd:
The `stdoutfd` field is set to the device file descriptor for `stdout` if the device is expecting to claim the `stdout` stream, or to the enumeration value `dev_not_claimed` otherwise.

stderrfd:
The `stderrfd` field is set to the device file descriptor for `stderr` if the device is expecting to claim the `stderr` stream, or to the enumeration value `dev_not_claimed` otherwise.

Registering New Devices

A new device can be registered with the following function:

```
int add_devtab_entry(DevEntry_t entry);
```

If the device is successfully registered, the `init()` routine of the device is called, with `entry` as its parameter. The `add_devtab_entry()` function returns the `DeviceID` of the device registered.

If the device is not successfully registered, a negative value is returned. Reasons for failure include (but are not limited to):

- The `DeviceID` is the same as another device, already registered
- There are no more slots left in the device registry table
- The `DeviceID` is less than zero
- Some of the function pointers are NULL
- The device's `init()` routine returned a failure result
- The device has attempted to claim a standard stream that is already claimed by another device

Pre-Registering Devices

The library source file `devtab.c`, which can be found under a VisualDSP installation in the subdirectory `...Blackfin\lib\src\libio`, declares the array:

```
DevEntry_t DevDrvTable[];
```

This array contains pointers to `DevEntry` structures for each device that is pre-registered, that is, devices that are available as soon as `main()` is entered, and that do not need to be registered at run-time by calling `add_devtab_entry()`. By default, the "*PrimIO*" device is registered. The `PrimIO` device provides support for target/host communication when using the simulators and the Analog Devices emulators and debug agents. This device is pre-registered, so that `printf()` and similar functions operate as expected without additional setup.

Additional devices can be pre-registered by the following process:

1. Take a copy of the `devtab.c` source file and add it to your project.
2. Declare your new device's `DevEntry` structure within the `devtab.c` file, for example,

   ```
   extern DevEntry myDevice;
   ```

3. Include the address of the `DevEntry` structure within the `DevDrvTable[]` array. Ensure that the table is null-terminated. For example,

   ```
   DevEntry_t DevDrvTable[MAXDEV] = {
   #ifdef PRIMIO
      &primio_deventry,
   #endif
      &myDevice,   /* new pre-registered device */
      0,
   };
   ```

All pre-registered devices are initialized automatically before `main()` is invoked, by the startup code. The run-time library calls the `init()` function of each of the pre-registered devices in turn.

The normal behavior of the `PrimIO` device when it is registered is to claim the first three files as `stdin`, `stdout` and `stderr`. These standard streams may be re-opened on other devices at run-time by using `freopen()` to close the `PrimIO`-based streams and reopen the streams on the current default device.

To allow an alternative device (either pre-registered or registered by `add_devtab_entry()`) to claim one or all of the standard streams:

1. Take a copy of the `primiolib.c` source file, and add it to your project.
2. Edit the appropriate `stdinfd`, `stdoutfd`, and `stderrfd` file descriptors in the `primio_deventry` structure to have the value `dev_not_claimed`.
3. Ensure the alternative device's `DevEntry` structure has set the standard stream file descriptors appropriately.

Both the device initialization routines, called from the startup code and `add_devtab_entry()`, return with an error if a device attempts to claim a standard stream that is already claimed.

Default Device

Once a device is registered, it can be made the default device using the following function:

```
void set_default_io_device(int);
```

The function should be passed the `DeviceID` of the device. There is a corresponding function for retrieving the current default device:

```
int get_default_io_device(void);
```

The default device is used by `fopen()` when a file is first opened. The `fopen()` function passes the open request to the `open()` function of the device indicated by `get_default_io_device()`. The device's file identifier (`fd`) returned by the `open()` function is private to the device; other devices may simultaneously have other open files that use the same identifier. An open file is uniquely identified by the combination of `DeviceID` and `fd`.

The `fopen()` function records the `DeviceID` and `fd` in the global open file table, and allocates its own internal `fid` to this combination. All future operations on the file use this `fid` to retrieve the `DeviceID` and thus direct the request to the appropriate device's primitive functions, passing the `fd` along with other parameters. Once a file has been opened by `fopen()`, the current value of `get_default_io_device()` is irrelevant to that file.

Remove and Rename Functions

The `PrimIO` device provides support for the `remove()` and `rename()` functions. These functions are not currently part of the extensible File I/O interface, since they deal purely with path names, and not with file descriptors. All calls to `remove()` and `rename()` in the run-time library are passed directly to the `PrimIO` device.

Default Device Driver Interface

The stdio functions provide access to the files on a host system through a device driver that supports a set of low-level I/O primitives. These low-level primitives are described under"Extending I/O Support To New Devices" on page 3-35. The default device driver implements these primitives based on a simple interface provided by the VisualDSP++ simulator and EZ-KIT Lite systems.

All the I/O requests submitted through the default device driver are channeled through the C function `_primIO`. The assembly label has two underscores, `__primIO` . The source for this function, and all the other library routines, can be found under the base installation for VisualDSP++ in the subdirectory `...Blackfin\lib\src\libio`.

The __primIO function accepts no arguments. Instead, it examines the I/O control block at the label _PrimIOCB. Without external intervention by a host environment, the __primIO routine simply returns, which indicates failure of the request. Two schemes for host interception of I/O requests are provided.

The first scheme is to modify control flow into and out of the __primIO routine. Typically, this would be achieved by a break point mechanism available to a debugger/simulator. Upon entry to __primIO, the data for the request resides in a control block at the label _PrimIOCB. If this scheme is used, the host should arrange to intercept control when it enters the __primIO routine, and, after servicing the request, return control to the calling routine.

The second scheme involves communicating with the DSP process through a pair of simple semaphores. This scheme is most suitable for an externally-hosted development board. Under this scheme, the host system should clear the data word whose label is __lone_SHARC; this causes __primIO to assume that a host environment is present and able to communicate with the process.

If __primIO sees that __lone_SHARC is cleared, then upon entry (for example, when an I/O request is made) it sets a nonzero value into the word labeled __Godot. The __primIO routine then busy-waits until this word is reset to zero by the host. The non-zero value of __Godot raised by __primIO is the address of the I/O control block.

Data Packing For Primitive I/O

The implementation of the __primIO interface is based on a word-addressable machine, with each word comprising a fixed number of 8-bit bytes. All READ and WRITE requests specify a move of some number of 8-bit bytes, that is, the relevant fields count 8-bit bytes, not words. Packing is always little endian, the first byte of a file read or written is the low-order byte of the first word transferred.

Data packing is set to one byte per word for the Blackfin processors. Data packing can be changed to accommodate other architectures by modifying the constant `BITS_PER_WORD`, defined in `_wordsize.h`. (For example, a processor with 16-bit addressable words would change this value to 16).

Note that the file name provided in an `OPEN` request uses the processor's "native" string format, normally one byte per word. Data packing applies only to `READ` and `WRITE` requests.

Data Structure for Primitive I/O

The I/O control block is declared in `_primio.h`, as follows.

```
typedef struct
{
   enum
   {
      PRIM_OPEN = 100,
      PRIM_READ,
      PRIM_WRITE,
      PRIM_CLOSE,
      PRIM_SEEK,
      PRIM_REMOVE,
      PRIM_RENAME
   } op;
   int   fileID;
   int   flags;
   unsigned char *buf;   /* data buffer, or file name     */
   int   nDesired;       /* number of characters to read  */
                         /* or write                      */
   int   nCompleted;     /* number of characters actually */
                         /* read or written               */
   void  *more;          /* for future use                */
}
PrimIOCB_T;
```

The first field, `op`, identifies which of the seven currently-supported operations is being requested.

The file ID for an open file is a non-negative integer assigned by the debugger or other "host" mechanism. The `fileID` values 0, 1, and 2 are pre-assigned to `stdin`, `stdout`, and `stderr`, respectively. No open request is required for these file IDs.

Before "activating" the debugger or other host environment, an OPEN or REMOVE request may set the `fileID` field to the length of the filename to open or delete; a RENAME request may also set the field to the length of the old filename. If the `fileID` field does contain a string length, then this will be indicated in the `flags` field (see below), and the debugger or other host environment will be able to use the information to perform a batch memory read to extract the filename. If the information is not provided, then the file name has to be extracted one character at a time.

The `flags` field is a bit field containing other information for special requests. Meaningful bit values for an `OPEN` operation are:

```
M_OPENR = 0x0001      /* open for reading                    */
M_OPENW = 0x0002      /* open for writing                    */
M_OPENA = 0x0004      /* open for append                     */
M_TRUNCATE = 0x0008   /* truncate to zero length if file exists */
M_CREATE = 0x0010     /* create the file if necessary        */
M_BINARY = 0x0020     /* binary file (vs. text file)         */
M_STRLEN_PROVIDED = 0x8000  /* length of file name(s) available */
```

For a `READ` operation, the low-order four bits of the `flag` value contain the number of bytes packed into each word of the read buffer, and the rest of the value is reserved for future use.

For a `WRITE` operation, the low-order four bits of the `flag` value contain the number of bytes packed into each word of the write buffer, and the rest of the value form a bit field, for which only the following bit is currently defined:

```
M_ALIGN_BUFFER = 0x10
```

If this bit is set for a WRITE request, the WRITE operation is expected to be aligned on a processor word boundary by writing padding NULs to the file before the buffer contents are transferred.

For an OPEN, REMOVE, and RENAME operation, the debugger (or other host mechanism) has to extract the filename(s) one character at a time from the memory of the target. However, if the bit corresponding to the value M_STRLEN_PROVIDED is set, then the I/O control block contains the length of the filename(s) and the debugger is able to use this information to perform a batch read of the target memory (see the description of the fields fileID and nCompleted).

For a SEEK request, the flags field indicates the seek mode (whence) as follows:

```
enum
{
   M_SEEK_SET = 0x0001,   /* seek origin is the start of
                             the file */
   M_SEEK_CUR = 0x0002,   /* seek origin is the current
                             position within the file */
   M_SEEK_END = 0x0004,   /* seek origin is the end of
                             the file */
};
```

The flags field is unused for a CLOSE request.

The buf field contains a pointer to the file name for an OPEN or REMOVE request, or a pointer to the data buffer for a READ or WRITE request. For a RENAME operation, this field contains a pointer to the old file name.

The nDesired field is set to the number of bytes that should be transferred for a READ or WRITE request. This field is also used by a RENAME request, and is set to a pointer to the new file name.

For a SEEK request, the nDesired field contains the offset at which the file should be positioned, relative to the origin specified by the flags field. (On architectures that only support 16-bit ints, the 32-bit offset at which the file should be positioned is stored in the combined fields [buf, nDesired]).

The nCompleted field is set by __primIO to the number of bytes actually transferred by a READ or WRITE operation. For a SEEK operation, __primIO sets this field to the new value of the file pointer. (On architectures that only support 16-bit ints, __primIO sets the new value of the file pointer in the combined fields [nCompleted, more]).

The RENAME operation may also make use of the nCompleted field. If the operation can determine the lengths of the old and new filenames, then it should store these sizes in the fields fileID and nCompleted, respectively, and also set the bit field flags to M_STRLEN_PROVIDED. The debugger (or other host mechanism) can then use this information to perform a batch read of the target memory to extract the file names. If this information is not provided, then each character of the file names will have to be read individually.

The more field is reserved for future use and currently is always set to NULL before calling _primIO.

Documented Library Functions

The C run-time library has several categories of functions and macros defined by the ANSI C standard, plus extensions provided by Analog Devices.

The following tables list the library functions documented in this chapter. Note that the tables list the functions for each header file separately; however, the reference pages for these library functions present the functions in alphabetical order.

Table 3-7. Library Functions in the `ctype.h` Header File

isalnum	isalpha	iscntrl
isdigit	isgraph	islower
isprint	ispunct	isspace
isupper	isxdigit	tolower
toupper		

Table 3-8. Library Functions in the `math.h` Header File

acos	asin	atan
atan2	ceil	cos
cosh	exp	fabs
floor	fmod	frexp
isinf	isnan	ldexp
log	log10	modf
pow	sin	sinh
sqrt	tan	tanh

Table 3-9. Library Functions in the `setjmp.h` Header File

longjmp	setjmp

Table 3-10. Library Functions in the `signal.h` Header File

raise	signal	interrupt

Table 3-11. Library Functions in the `stdarg.h` Header File

va_arg	va_end	va_start

Table 3-12. Supported Library Functions in the `stdio.h` Header File

clearerr	fclose	feof
ferror	fflush	fgetc
fgetpos	fgets	fprintf
fputc	fputs	fopen
fread	freopen	fscanf
fseek	fsetpos	ftell
fwrite	getc	getchar
gets	perror	printf
putc	putchar	puts
remove	rename	rewind
scanf	setbuf	setvbuf
sprintf	sscanf	ungetc
vfprintf	vprintf	vsprintf

Table 3-13. Library Functions in `stdlib.h` Header File

abort	abs	atexit
atof	atoi	atol
atold	atoll	
bsearch	calloc	div
exit	free	heap_calloc
heap_free	heap_init	heap_install
heap_lookup	heap_malloc	heap_realloc
heap_space_unused	labs	ldiv
malloc	qsort	rand
realloc	space_unused	srand
strtod	strtof	strtol
strtold	strtoll	strtoul
strtoull		

Table 3-14. Library Functions in `string.h` Header File

memchr	memcmp	memcpy
memmove	memset	strcat
strchr	strcmp	strcoll
strcpy	strcspn	strerror
strlen	strncat	strncmp
strncpy	strpbrk	strrchr
strspn	strstr	strtok
strxfrm		

Table 3-15. Library Functions in `time.h` Header File

asctime	clock	ctime
difftime	gmtime	localtime
mktime	strftime	time

C Run-Time Library Reference

The C run-time library is a collection of functions called from your C programs. Note the following items apply to all of the functions in the library.

Notation Conventions
An interval of numbers is indicated by the minimum and maximum, separated by a comma, and enclosed in two square brackets, two parentheses, or one of each. A square bracket indicates that the endpoint is included in the set of numbers; a parenthesis indicates that the endpoint is not included.

Reference Format
Each function in the library has a reference page. These pages have the following format:

Name and Purpose of the function

Synopsis – Required header file and functional prototype

Description – Function specification

Error Conditions – Method that the functions use to indicate an error

Example –Typical function usage

See Also – Related functions

abort

abnormal program end

Synopsis

```
#include <stdlib.h>
void abort(void);
```

Description

The `abort` function causes an abnormal program termination by raising the `SIGABRT` exception. If the `SIGABRT` handler returns, `abort()` calls `exit()` to terminate the program with a failure condition.

Error Conditions

The `abort` function does not return.

Example

```
#include <stdlib.h>
extern int errors;

if(errors)    /* terminate program if */
   abort();   /* errors are present   */
```

See Also

atexit, exit

abs

absolute value

Synopsis

```
#include <stdlib.h>
int abs(int j);
```

Description

The `abs` function returns the absolute value of its integer input.

Note: `abs(INT_MIN)` returns `INT_MIN`.

Error Conditions

The `abs` function does not return an error condition.

Example

```
#include <stdlib.h>
int i;
i = abs(-5);       /* i == 5 */
```

See Also

fabs, labs

acos

arc cosine

Synopsis

```
#include <math.h>
double acos (double x);
float acosf (float x);
long double acosd (long double x);
fract16 acos_fr16 (fract16 x);
```

Description

The acos functions return the arc cosine of `x`. Both the argument `x` and the function results are in radians.

The input for the functions `acos`, `acosf`, and `acosd` must be in the range [-1, 1], and the functions return a result that will be in the range [0, π].

The `acos_fr16` function is defined for fractional input values between 0 and 0.9. The output from the function is in the range [`acos(0)*2/`π, `acos(0.9)*2/`π].

Error Conditions

The acos functions return a zero if the input is not in the defined range.

Example

```
#include <math.h>
double y;
y = acos(0.0);       /* y = π/2 */
```

See Also

cos

adi_acquire_lock

Obtain and release locks for multi-core synchronization

Synopsis

```
#include <ccblkfn.h>
void adi_acquire_lock(testset_t *);
int adi_try_lock(testset_t *);
void adi_release_lock(testset_t *);
```

Description

These functions provide locking facilities for multi-core applications that need to ensure private access to shared resources, or for applications that need to build synchronization mechanisms.

The functions operate on a pointer to a `testset_t` object, which is a private type used only by these routines. Objects of type `testset_t` must be global, and initialized to zero (which indicates that the lock is unclaimed). The type is automatically volatile.

The `adi_acquire_lock` function repeatedly attempts to acquire the lock, until successful. Upon return, the lock will have been acquired. The function does not make use of any timers or other mechanisms to pause between attempts, so this functionimplies continuous accesses to the lock object.

The `adi_try_lock` function makes a single attempt to acquire the lock, but does not block if the lock has already been acquired. The function returns non-zero if it has successfully acquired the lock, and zero if the lock was not available.

The `adi_release_lock` function releases the lock object, marking it as available to the next attempt by `adi_acquire_lock` or `adi_try_lock`. The `adi_release_lock` function does not return a value, and does not verify whether the caller already holds the lock, or even if the lock is already held by “another” caller.

Error Conditions

These functions do not return error conditions. Neither `adi_acquire_lock()` nor `adi_release_lock()` return values. The `adi_try_lock()` function merely returns a value indicating whether the lock was acquired.

Examples

```
#include <ccblkfn.h>

void add_one(testset_t *lockptr, volatile int *valptr)
{
      adi_acquire_lock(lockptr);   // blocks
      valptr += 1;
      adi_release_lock(lockptr);
}
```

```
#include <ccblkfn.h>

void claim_lock(testset_t *lockptr, int how_many_millisecs)
{
      unsigned int intr_mask
      intr_mask = cli();
      while (!adi_try_lock(lockptr)) {
            sti(intr_mask);
            sleep_for_a_while(how_many_millisecs);
            intr_mask = cli();
      }
      sti(intr_mask);
}
```

To be useful, the `testset_t` object must be located in a shared area of memory accessible by both cores.
These functions do not disable interrupts; that is the responsibility of the caller.

See Also

adi_core_id, adi_obtain_mc_slot

adi_core_id

Identify caller's core

Synopsis

```
#include <ccblkfn.h>
int adi_core_id(void);
```

Description

The `adi_core_id` function returns a numeric value indicating which processor core is executing the call to the function. This function is most useful on multi-core processors, when the caller is a function shared between both cores, but which needs to do different actions (or access different data) depending on the core executing it.

The function returns a zero value when executed by Core A, and a value of one when executed on Core B.

Error Conditions

The `adi_core_id` function does not return an error condition

Example

```
#include <ccblkfn.h>

const char *core_name(void)
{
      if (adi_core_id() == 0)
         return "Core A";
   else
         return "Core B";
}
```

See Also

adi_acquire_lock, adi_obtain_mc_slot

adi_obtain_mc_slot

Obtain and manage storage for multi-core private data in shared functions.

Synopsis

```
#include <mc_data.h>
adi_obtain_mc_slot(int *slotID, void (*fn)(void *));
int adi_free_mc_slot(int slotID);
int adi_set_mv_value(int slotID, void *valptr);
void *adi_get_mc_value(int slotID);
```

Description

These functions provide a framework for shared functions that may be called from any core in a multi-core environment, yet need to maintain data values that are private to the calling core. An example is `errno`—in a multi-core environment, each core needs to maintain its own version of the errno value, but the correct version of `errno` must be updated when a shared Standard library function is called.

The framework operates by maintaining a set of "slots", each slot corresponds to a data object that must be core-local. The slot holds a pointer for each core, which can be set to point to the core's private version of the data object.

The process is as follows:

1. If this is the first time any core has needed the private data, then allocate a slot.
2. If this is the first time this core has needed the private data, then allocate storage for the data and record it in the slot, else retrieve the location of the data's storage from the slot.
3. Access the data.

The adi_obtain_mc_slot function is called to allocate a slot, when no core has previously needed to access the data. slotID must be a pointer to a global variable, shared by all the cores, which is initialized to the value adi_mc_unallocated. The fn parameter must be NULL.

If the adi_obtain_mc_slot function can allocate a slot for the data object, it will return the slot's identifier, via the slotID pointer, and will return a non-zero value. If there are no more slots remaining, the function returns a zero value.

The adi_free_mc_slot function releases the slot indicated by slotID, which must have been previously allocated by the adi_obtain_mc_slot function. If slotID indicate a valid slot, the slot is freed and the function returns a non-zero value. If slotID does not indicate a currently-valid slot, the function returns zero.

The adi_set_mc_value function records the valptr pointer in the slot indicated by slotID, as the location of the private data object for the calling core. The function returns true if slotID refers to a currently-valid slot, otherwise the function returns false.

The adi_get_mc_value function returns a pointer previously stored in the slot indicated by slotID, for the calling core. The pointer must have been previously stored by the adi_set_mc_value function, by the current core, otherwise the function returns NULL. The function also returns NULL if slotID does not indicate a currently-valid slot

Error Conditions

The adi_obtain_mc_slot function returns a zero value if a new slot cannot be allocated.

The adi_free_mc_slot and adi_set_mc_value functions both return a zero value if slotID does not refer to a currently-valid slot.

The `adi_get_mc_value` function returns NULL if `slotID` does not refer to a currently-valid slot, or if the calling core has not yet stored a pointer in the slot via `adi_set_mc_value`.

Examples

```
example: error handling omitted */
#include <mc_data.h.h>
#include <ccblkfn.h>
#include <stdlib.h>

static int slotid = adi_mc_unallocated);
static testset_t slotlock = 0;

void set_error(int val)
{
      int *storage;
      adi_acquire_lock(&slotlock);
      if (slotid == adi_mc_unallocated {
            // first core here
            adi_obtain_mc_slot(&slotid, NULL);
      }
      adi_release_lock(&slotlock);
      storage = adi_get_mc_value(slotid);
      if (storage == NULL) {
            // first time this core is here
            storage = malloc(sizeof(int));
            adi_set_mc_value(slotid, storage);
      }
      *storage = val;
}
```

The multi-core private storage routines do not disable interrupts; that is left at the caller's discretion.

asctime

convert broken-down time into a string

Synopsis

```
#include <time.h>
char *asctime(const struct tm *t);
```

Description

The `asctime` function converts a broken-down time, as generated by the functions `gmtime` and `localtime`, into an ASCII string that will contain the date and time in the form

```
DDD MMM dd hh:mm:ss YYYY\n
```

where

- `DDD` represents the day of the week (that is. Mon, Tue, Wed, etc.)
- `MMM` is the month and will be of the form Jan, Feb, Mar, etc
- `dd` is the day of the month, from 1 to 31
- `hh` is the number of hours after midnight, from 0 to 23
- `mm` is the minute of the day, from 0 to 59
- `ss` is the second of the day, from 0 to 61 (to allow for leap seconds)
- `YYYY` represents the year

The function returns a pointer to the ASCII string, which may be over-written by a subsequent call to this function. Also note that the function `ctime` returns a string that is identical to

```
asctime(localtime(&t))
```

Error Conditions

The `asctime` function does not return an error condition.

Example

```
#include <time.h>
#include <stdio.h>

struct tm tm_date;

printf("The date is %s",asctime(&tm_date));
```

See Also

ctime, gmtime, localtime

asin

arc sine

Synopsis

```
#include <math.h>

double asin (double x);
float asinf (float x);
long double asind (long double x);

fract16 asin_fr16(fract16 x);
```

Description

The asin functions return the arc sine of the argument `x`. Both the argument `x` and the function results are in radians.

The input for the functions `asin`, `asinf`, and `asind` must be in the range [-1, 1], and the functions return a result that will be the range [-π/2, π/2].

The `asin_fr16` function is defined for fractional input values in the range [-0.9, 0.9]. The output from the function is in the range [`asin(-0.9)*2/`π, `asin(0.9)*2/`π].

Error Conditions

The asin functions return a zero if the input is not in the defined range.

Example

```
#include <math.h>
double y;
y = asin(1.0);       /* y = π/2 */
```

See Also

sin

atan

arc tangent

Synopsis

```
#include <math.h>

double atan (double x);
float atanf (float x);
long double atand (long double x);

fract16 atan_fr16 (fract16 x);
```

Description

The atan functions return the arc tangent of the argument. Both the argument `x` and the function results are in radians.

The `atanf`, `atan`, and `atand` functions return a result that is in the range [-π/2, π/2].

The `atan_fr16` function is defined for fractional input values in the range [-1.0, 1.0). The output from the function is in the range [-π/4, π/4].

Error Conditions

The atan functions do not return an error condition.

Example

```
#include <math.h>
double y;
y = atan(0.0);      /* y = 0.0 */
```

See Also

atan2, tan

atan2

arc tangent of quotient

Synopsis

```
#include <math.h>

double atan2 (double x, double y);
float atan2f (float x, float y);
long double atan2d (long double x, long double y);

fract16 atan2_fr16 (fract16 x, fract16 y);
```

Description

The atan2 functions compute the arc tangent of the input value `y` divided by input value `x`. The output is in radians.

The `atan2f`, `atan2`, and `atan2d` functions return a result that is in the range [-π, π].

The `atan2_fr16` function is defined for fractional input values in the range [-1.0, 1.0). The output from this function is scaled by π and is in the range [-1.0, 1.0).

Error Conditions

The atan2 functions return a zero if `x=0` and `y <> 0`.

Example

```
#include <math.h>
double a;
float b;

a = atan2 (0.0, 0.5);     /* the error condition: a = 0.0 */
b = atan2f (1.0, 0.0);    /* b = π/2                      */
```

See Also

atan, tan

atexit

register a function to call at program termination

Synopsis

```
#include <stdlib.h>
int atexit(void (*func)(void));
```

Description

The `atexit` function registers a function to be called at program termination. Functions are called once for each time they are registered, in the reverse order of registration. Up to 32 functions can be registered using the `atexit` function.

Error Conditions

The `atexit` function returns a nonzero value if the function cannot be registered.

Example

```
#include <stdlib.h>
extern void goodbye(void);

if (atexit(goodbye))
   exit(1);
```

See Also

abort, exit

atof

convert string to a double

Synopsis

```
#include <stdlib.h>
double atof(const char *nptr);
```

Description

The `atof` function converts a character string into a floating-point value of type `double`, and returns its value. The character string is pointed to by the argument `nptr` and may contain any number of leading whitespace characters (as determined by the function `isspace`) followed by a floating-point number. The floating-point number may either be a decimal floating-point number or a hexadecimal floating-point number.

A decimal floating-point number has the form:

```
[sign] [digits] [.digits] [{e|E} [sign] [digits]]
```

The `sign` token is optional and is either plus (+) or minus (–); and `digits` are one or more decimal digits. The sequence of digits may contain a decimal point (.).

The decimal digits can be followed by an exponent, which consists of an introductory letter (`e` or `E`) and an optionally signed integer. If neither an exponent part nor a decimal point appears, a decimal point is assumed to follow the last digit in the string.

The form of a hexadecimal floating-point number is:

```
[sign] [{0x}|{0X}] [hexdigs] [.hexdigs] [{p|P} [sign] [digits]]
```

A hexadecimal floating-point number may start with an optional plus (+) or minus (–) followed by the hexadecimal prefix `0x` or `0X`. This character sequence must be followed by one or more hexadecimal characters that optionally contain a decimal point (.).

The hexadecimal digits are followed by a binary exponent that consists of the letter `p` or `P`, an optional sign, and a non-empty sequence of decimal digits. The exponent is interpreted as a power of two that is used to scale the fraction represented by the tokens `[hexdigs] [.hexdigs]`.

The first character that does not fit either form of number stops the scan.

Error Conditions

The `atof` function returns a zero if no conversion is made. If the correct value results in an overflow, a positive or negative (as appropriate) `HUGE_VAL` is returned. If the correct value results in an underflow, 0.0 is returned. The `ERANGE` value is stored in `errno` in the case of either an overflow or underflow.

Notes

The function reference `atof (pdata)` is functionally equivalent to:

```
strtod (pdata, (char *) NULL);
```

and therefore, if the function returns zero, it is not possible to determine whether the character string contained a (valid) representation of 0.0 or some invalid numerical string.

Example

```
#include <stdlib.h>
double x;

x = atof("5.5");      /* x == 5.5 */
```

See Also

atoi, atol, strtod

atoi

convert string to integer

Synopsis

```
#include <stdlib.h>
int atoi (const char *nptr);
```

Description

The `atoi` function converts a character string to an integer value. The character string to be converted is pointed to by the input pointer, `nptr`. The function clears any leading characters for which `isspace` would return true. Conversion begins at the first digit (with an optional preceding sign) and terminates at the first non-digit.

Error Conditions

The `atoi` function returns a zero if no conversion is made.

Example

```
#include <stdlib.h>
int i;

i = atoi("5");       /* i == 5 */
```

See Also

atof, atol, strtod, strtol, strtoul

atol

convert string to long integer

Synopsis

```
#include <stdlib.h>
long atol (const char *nptr);
```

Description

The `atol` function converts a character string to a `long` integer value. The character string to be converted is pointed to by the input pointer, `nptr`. The function clears any leading characters for which `isspace` would return true. Conversion begins at the first digit (with an optional preceding sign) and terminates at the first non-digit.

There is no way to determine if a zero is a valid result or an indicator of an invalid string.

Error Conditions

The `atol` function returns a zero if no conversion is made.

Example

```
#include <stdlib.h>
long int i;

i = atol("5");       /* i == 5 */
```

See Also

atof, strtod, strtol, strtoul

atold

convert string to a long double

Synopsis

```
#include <stdlib.h>
long double atold(const char *nptr);
```

Description

The `atold` function converts a character string into a floating-point value of type `long double`, and returns its value. The character string is pointed to by the argument `nptr` and may contain any number of leading whitespace characters (as determined by the function `isspace`) followed by a floating-point number. The floating-point number may either be of the form of a decimal floating-point number or a hexadecimal floating-point number.

A decimal floating-point number has the form:

```
[sign] [digits] [.digits] [{e|E} [sign] [digits]]
```

The `sign` token is optional and is either plus (+) or minus (–); and `digits` are one or more decimal digits. The sequence of digits may contain a decimal point (.).

The decimal digits can be followed by an exponent, which consists of an introductory letter (`e` or `E`) and an optionally signed integer. If neither an exponent part nor a decimal point appears, a decimal point is assumed to follow the last digit in the string.

The form of a hexadecimal floating-point number is:

```
[sign] [{0x}|{0X}] [hexdigs] [.hexdigs] [{p|P} [sign] [digits]]
```

A hexadecimal floating-point number may start with an optional plus (+) or minus (-) followed by the hexadecimal prefix `0x` or `0X` . This character sequence must be followed by one or more hexadecimal characters that optionally contain a decimal point (.).

The hexadecimal digits are followed by a binary exponent that consists of the letter `p` or `P` , an optional sign, and a non-empty sequence of decimal digits. The exponent is interpreted as a power of two that is used to scale the fraction represented by the tokens `[hexdigs][.hexdigs]`.

The first character that does not fit either form of number stops the scan.

Error Conditions

The `atold` function returns a zero if no conversion could be made. If the correct value results in an overflow, a positive or negative (as appropriate) `LDBL_MAX` is returned. If the correct value results in an underflow, 0.0 is returned. The `ERANGE` value is stored in `errno` in the case of either an overflow or underflow.

Notes

The function reference `atold (pdata)` is functionally equivalent to:

```
strtold (pdata, (char *) NULL);
```

and therefore, if the function returns zero, it is not possible to determine whether the character string contained a (valid) representation of 0.0 or some invalid numerical string.

Example

```
#include <stdlib.h>
long double x;

x = atold("5.5");       /* x == 5.5 */
```

See Also

atoi, atol, strtold

atoll

convert string to long long integer

Synopsis

```
#include <stdlib.h>
long long atoll (const char *nptr);
```

Description

The `atoll` function converts a character string to a `long long` integer value. The character string to be converted is pointed to by the input pointer, `nptr`. The function clears any leading characters for which `isspace` would return true. Conversion begins at the first digit (with an optional preceding sign) and terminates at the first non-digit.

There is no way to determine if a zero is a valid result or an indicator of an invalid string.

Error Conditions

The `atoll` function returns a zero if no conversion is made.

Example

```
#include <stdlib.h>
long long int i;

i = atoll("5");      /* i == 5 */
```

See Also

strtoll

bsearch

perform binary search in a sorted array

Synopsis

```
#include <stdlib.h>
void *bsearch (const void *key, const void *base,
               size_t nelem, size_t size,
               int (*compare)(const void *, const void *));
```

Description

The `bsearch` function executes a binary search operation on a presorted array, where:

- `key` points to the element to search for
- `base` points to the start of the array
- `nelem` is the number of elements in the array
- `size` is the size of each element of the array
- `*compare` points to the function used to compare two elements. It takes two parameters—a pointer to the key, and a pointer to an array element. It should return a value less than, equal to, or greater than zero, according to whether the first parameter is less than, equal to, or greater than the second.

The `bsearch` function returns a pointer to the first occurrence of `key` in the array.

Error Conditions

The `bsearch` function returns a null pointer when the key is not found in the array.

Example

```
#include <stdlib.h>
char *answer;
char base[50][3];

answer = bsearch("g", base, 50, 3, strcmp);
```

See Also

qsort

cache_invalidate

invalidate processor instruction and data caches

Synopsis

```
#include <cplbtab.h>
void cache_invalidate(int cachemask);
void icache_invalidate(void);
void dcache_invalidate(int a_or_b);
void dcache_invalidate_both(void);
```

Description

The `cache_invalidate` function and its related functions `icache_invalidate` and `dcache_invalidate` invalidate the contents of the processor's instruction and data caches, forcing any data to be re-fetched from memory. Modified data cached in Write Back mode is **not** flushed to memory first.

The `cache_invalidate` routine calls its support routines according to the bits set in parameter `cachemask`. The bits have the following meanings.

Bit set	Meaning
`CPLB_ENABLE_ICACHE`	Invalidate instruction cache
`CPLB_ENABLE_DCACHE`	Invalidate data cache A
`CPLB_ENABLE_DCACHE2`	Invalidate data cache B

A call is made to the appropriate support routine for each bit set. If bits are set to indicate that both data cache A and data cache B must be invalidated, a single call is made to the `dcache_invalidate_both` routine.

On the ADSP-BF535 and AD6532 processors, `cache_invalidate` is called by the default start-up code on reset, and is passed the value of `___cplb_ctrl` as its parameter. Thus, each enabled cache is invalidated during start-up. On other Blackfin processors, the caches automatically reset to the "invalidated" state, and no call is necessary, nor performed.

The `dcache_invalidate` routine only invalidates a single data cache, selected by its `a_or_b` parameter:

a_or_b Value	Meaning
`CPLB_INVALIDATE_A`	Invalidate data cache A
`CPLB_INVALIDATE_B`	Invalidate data cache B

The `dcache_invalidate_both` routine invalidates both data cache A and data cache B. On the ADSP-BF535 and AD6532 processors, it is implemented by calling `dcache_invalidate` for each data cache in turn. On other Blackfin processors, the routine toggles the bits of the `DMEM_CONTROL` register to invalidate all contents of both data caches at once, and is considerably faster than calling `dcache_invalidate` for each data cache separately.

Error Conditions

The cache invalidation routines do not return an error condition.

Example

```
#include <cplbtab.h>

void clean_cache(int which)
{
   switch (which) {
   case 1:
         icache_invalidate();
         break;
   case 2:
```

```
            dcache_invalidate(CPLB_INVALIDATE_A);
            break;
        case 4:
            dcache_invalidate(CPLB_INVALIDATE_B);
            break;
        case 6:
            dcache_invalidate_both();
            break;
        default:
            cache_invalidate(__cplb_ctrl);
            break;
        }
    }
```

See Also

flush_data_cache

calloc

allocate and initialize memory

Synopsis

```
#include <stdlib.h>
void *calloc (size_t nmemb, size_t size);
```

Description

The `calloc` function dynamically allocates a range of memory and initializes all locations to zero. The number of elements (the first argument) multiplied by the size of each element (the second argument) is the total memory allocated. The memory may be deallocated with the `free` function.

Error Conditions

The `calloc` function returns a null pointer if unable to allocate the requested memory.

Example

```
#include <stdlib.h>
int *ptr;

ptr = (int *) calloc(10, sizeof(int));
      /* ptr points to a zeroed array of length 10 */
```

See Also

free, malloc, realloc

ceil

ceiling

Synopsis

```
#include <math.h>

double ceil (double x);
float ceilf (float x);
long double ceild (long double x);
```

Description

The ceil functions return the smallest integral value that is not less than the argument `x`.

Error Conditions

The ceil functions do not return an error condition.

Example

```
#include <math.h>
double y;
float x;

y = ceil (1.05);       /* y = 2.0  */
x = ceilf (-1.05);     /* y = -1.0 */
```

See Also

floor

clock

processor time

Synopsis

```
#include <time.h>
clock_t clock(void);
```

Description

The `clock` function returns the number of processor cycles that have elapsed since an arbitrary starting point. The function returns the value `(clock_t) -1`, if the processor time is not available or if it cannot be represented. The result returned by the function may be used to calculate the processor time in seconds by dividing it by the macro `CLOCKS_PER_SEC`. For more information, see "time.h" on page 3-24. An alternative method of measuring the performance of an application is described in "Measuring Cycle Counts" on page 4-33.

Error Conditions

The `clock` function does not return an error condition.

Example

```
#include <time.h>

time_t start_time,stop_time;
double time_used;

start_time = clock();
compute();
stop_time = clock();

time_used = ((double) (stop_time - start_time)) / CLOCKS_PER_SEC;
```

See Also

No references to this function.

cos

cosine

Synopsis

```
#include <math.h>

double cos (double x);
float cosf (float x);
long double cosd (long double x);

fract16 cos_fr16 (fract16 x);
```

Description

The cos functions return the cosine of the argument. Both the argument `x` and the results returned by the functions are in radians.

The result returned by the functions `cos`, `cosf`, and `cosd` is in the range [-1, 1].

The `cos_fr16` function inputs a fractional value in the range [-1.0, 1.0) corresponding to [-π/2, π/2]. The domain represents half a cycle which can be used to derive a full cycle if required (see Notes below). The result, in radians, is in the range [-1.0, 1.0).

Error Conditions

The cos functions do not return an error condition.

Example

```
#include <math.h>
double y;
y = cos(4.04159);      /* y = -1.0 */
```

Notes

The domain of the `cos_fr16` function is restricted to the fractional range [`0x8000`, `0x7fff`] which corresponds to half a period from –(π /2) to π/2. It is possible however to derive the full period using the following properties of the function.

```
cosine [0, π/2] = -cosine [π, 3/2 π]
cosine [-π/2, 0] = -cosine [π/2, π]
```

The function below uses these properties to calculate the full period (from 0 to 2π) of the `cosine` function using an input domain of `[0, 0x7fff]`.

```
#include <math.h>

fract16 cos2pi_fr16 (fract16 x)
{
   if (x < 0x2000) {                    /* <0.25  */
      /* first quadrant [0..π/2):                */
      /* cos_fr16([0x0..0x7fff]) = [0..0x7fff)   */
      return cos_fr16(x * 4);

   } else if (x < 0x6000) {             /* < 0.75  */
      /* if (x < 0x4000)                         */
      /* second quadrant [π/2..π):               */
      /* -cos_fr16([0x8000..0x0)) = [0x7fff..0)  */
      /*                                         */
      /* if (x < 0x6000)                         */
      /* third quadrant [π..3/2π):               */
      /* -cos_fr16([0x0..0x7fff]) = [0..0x8000)  */
      return -cos_fr16((0xc000 + x) * 4);

   } else {
      /* fourth quadrant [3/2π..π):              */
      /* cos_fr16([0x8000..0x0)) = [0x8000..0)   */
      return cos_fr16((0x8000 + x) * 4);
   }
}
```

See Also

acos, sin

cosh

hyperbolic cosine

Synopsis

```
#include <math.h>

double cosh (double x);
float coshf (float x);
long double coshd (long double x);
```

Description

The cosh functions return the hyperbolic cosine of their argument.

Error Conditions

The domain of `coshf` is [-89.39, 89.39], and the domain for `coshd` is [-710.44, 710.44]. The functions return `HUGE_VAL` if the input argument `x` is outside the respective domains.

Example

```
#include <math.h>
double x, y;
float v, w;

y = cosh (x);
v = coshf (w);
```

See Also

sinh

cplb_hdr

default exception handler for memory-related events

Synopsis

```
#include <cplbtab.h>
void cplb_hdr (void);
```

Description

The `cplb_hdr` routine is the default exception handler, installed by the default start-up code to service CPLB-related events if CPLBs or caching is indicated by the `__cplb_ctrl` variable.

The routine saves the processor context, before examining the exception details to determine the kind of exception raised. If it is an Instruction CPLB Miss, a Data CPLB Miss, or a Data CPLB Write, then the routine invokes `cplb_mgr` to handle the event, otherwise it calls the routine `_unknown_exception_occurred`, which is not expected to return.

If `cplb_mgr` indicates a successful handling, the routine returns from the exception, restoring the context as it does. Otherwise, it invokes an appropriate diagnostic routine.

Error Conditions

The `cplb_hdr` routine calls other routines to deal with each of the error codes returned by the `cplb_mgr` routine. By default, these routines are stubs that loop forever—you can replace them with your own routines if you wish to provide more detailed handling.

The error codes are:

Error code	Response
CPLB_RELOADED	Indicates success; returns
CPLB_NO_ADDR_MATCH	Calls stub _cplb_miss_without_replacement
CPLB_NO_UNLOCKED	Calls stub _cplb_miss_all_locked
CPLB_PROT_VIOL	Calls stub _cplb_protection_violation
others	Loops at label strange_return_from_cplb_mgr

Example

```
#include <cplbtab.h>
#include <sys/exception.h>

void setup_cplb_handling(void) {
     register_handler(ik_exception, (ex_handler_fn)cplb_hdr);
}
```

See also `Blackfin\lib\src\libc\crt\basiccrt.s` in the VisualDSP++ installation directory.

See Also

cplb_init, cplb_mgr

cplb_init

initialize CPLBs and caches at start-up

Synopsis

```
void cplb_init(int bitmask);
```

Description

The `cplb_init` routine is called by the default start-up code during processor initialization. It initializes the memory protection hardware and enables caches where requested, according to configuration data in two tables. It is not expected that `cplb_init()` is called from normal user code, nor is it expected that it is called more than once following each processor reset.

The routine's behavior is controlled by the following data structures:

- the `__cplb_ctrl` variable
- the `dcplbs_table[]` array
- the `icplbs_table[]` array

Initially, the routine tests the `__cplb_ctrl` variable to determine whether any of the caches have been enabled when the `.LDF` file has already mapped code or data into the corresponding cache area. If so, this would lead to corrupted code or data; therefore, the `cplb_init` routine aborts by jumping to infinite loops labelled with diagnostic symbols, for example, `l1_code_cache_enabled_when_l1_used_for_code`.

For the ADSP-BF535 and AD6532 processors, if caches are indicated by the `__cplb_ctrl` variable, then the routine invokes the `cache_invalidate` routine to first invalidate the caches, so that they are not enabled while containing random bits.

For each of the data and instruction CPLBs, if requested, the `cplb_init` routine copies from one to sixteen entries from the configuration tables, installing the tables' entries into the corresponding registers. For example, `icplbs_table[0]` is copied into `ICPLB_DATA0` and `ICPLB_ADDR0`, and `dcplbs_table[0]` is copied into `DCPLB_DATA0` and `DCPLB_ADDR0`.

The copying is not verbatim; if caches are not requested by the `__cplb_ctrl` variable, cache bits are masked off the values written to the `xCPLB_DATA` *n* registers.

If a table has from one to sixteen entries, all of the table's entries are installed, with any unused `xCPLB_DATA` *n* registers being marked as "Invalid". If a table contains more entries, then only the first sixteen are installed. It is assumed that an appropriate exception handler was installed to process any CPLB Miss exceptions that occur. The `cplb_hdr` routine is an example of such an exception handler.

After installing the CPLB entries from the tables, the `cplb_init` routine modifies the `IMEM_CONTROL` and `DMEM_CONTROL` registers to enable the CPLBs and caches that were indicated. The `cplb_init` routine also sets the following `DMEM_CONTROL` bits:

- The Data Cache Bank Select bit is set, according to whether `CPLB_ENABLE_DCBS` is set.
- The `DAG0/1` Port Preference bits are set to 1 and 0, respectively, to reduce memory access stalls.

Error Conditions

The `cplb_init` routine does not return an error condition. If it encounters an error during initialization, it jumps to a label indicative of the problem.

Example

See the source for the start-up code, in the VisualDSP++ installation directory, under `...Blackfin\lib\src\libc\crt\basiccrt.s`.

See Also

cplb_hdr

cplb_mgr

CPLB management routine for CPLB exceptions.

Synopsis

```
#include <cplbtab.h>
int cplb_mgr(int event, int bitmask);
```

Description

The `cplb_mgr` routine manages the active CPLB tables for instructions and data. It is intended to be invoked from an exception handler upon receipt of a CPLB-related `event`. `cplb_hdr`, installed by the default start-up code, is a typical example of such a handler.

The event parameter indicates the action that the routine should take:

Event Value	Required Action
`CPLB_EVT_ICPLB_MISS`	Replace an active Instruction CPLB
`CPLB_EVT_DCPLB_MISS`	Replace an active Data CPLB
`CPLB_EVT_DCPLB_WRITE`	Mark an existing Data CPLB as Dirty

To replace an Instruction CPLB, the routine determines the faulting address from the processor's registers, and searches the `icplbs_table[]` looking for an entry whose start address and size addresses a region of memory that includes the faulting address. If none is found, the routine returns `CPLB_NO_ADDR_MATCH` to indicate that the faulting address is not covered by any of the entries in `icplbs_table[]`, and is therefore an invalid address.

The routine selects the first active Instruction CPLB that is not locked, and shuffles all following Instruction CPLBs up, overwriting it, so that the last Instruction CPLB is free to be used by the entry to be installed. In this

manner, the replacement algorithm is Least Recently Installed. If no unlocked Instruction CPLBs can be found, the routine returns `CPLB_NO_UNLOCKED`.

The new Instruction CPLB is installed into `ICPLB_ADDR15` and `ICPLB_DATA15`. If the bitmask parameter indicates that Instruction Cache is not enabled (that is, if bitmask does not have bit `CPLB_ENABLE_ICACHE` set), then cache bits are masked off the entry as it is installed.

Replacing a Data CPLB follows a similar process, but must also deal with CPLBs that indicate data is to be cached in Write Back mode. In this case, the routine first attempts to select a Clean Data CPLB to evict. If no unlocked Clean Data CPLB can be found, then the routine falls back on selecting Dirty Data CPLBs. As for Instruction CPLBs, if no unlocked CPLB can be selected, the routine returns `CPLB_NO_UNLOCKED`.

If it is necessary to evict a Dirty Data CPLB, the `cplb_mgr` routine first flushes any modified cache entries corresponding to the victim Data CPLB's memory page, forcing the modified data to be written back to secondary memory.

When the new Data CPLB is installed, cache bits are masked off if the bitmask parameter does not have either of the `CPLB_ENABLE_DCACHE` or `CPLB_ENABLE_DCACHE2` bits sets.

The `cplb_mgr` routine is called to handle Data CPLB Write events when a page is cached in Write Back mode, and the first write occurs to a Clean page. In this case, the routine locates the active CPLB using the processor's MMRs and verifies that it is a clean, cached, Write Back CPLB, and marks the page as Dirty. Future writes to the page do not trigger an exception, but now the page has been marked as Dirty, so the routine can flush the modified data from the cache if it becomes necessary to evict the page.

If the page indicated by a Data CPLB Write is not a Clean, cached Write Back page, this indicates that a protection violation has occurred, for example, a write to a Supervisor-only page while in User Mode, and, therefore, the routine returns `CPLB_PROT_VIOL`.

When a Data CPLB is installed during the eviction process, pages that are cached in Write Back mode are not forced to be marked as Clean. This is because there is a performance trade-off that the application designer can exploit if the expectation is that Data CPLB Miss exceptions are very rare. A Dirty Data CPLB is expensive to evict, because of the cost of flushing modified data to secondary memory, but if no eviction is ever expected, this is irrelevant. In contrast, if a page is expected to be modified but never flushed, a Clean Data CPLB will pay the cost of a Data CPLB Write exception on first write. Therefore, the designer may choose to pre-mark pages as Dirty, with the expectation that the CPLB eviction process will never occur.

Error Conditions

The `cplb_mgr` routine returns the following values:

Value	Meaning
`CPLB_RELOADED`	The event was serviced successfully
`CPLB_NO_ADDR_MATCH`	There is no entry in the appropriate `cplbs_table[]` that covers the faulting address.
`CPLB_NO_UNLOCKED`	All the active CPLBs are marked as “locked” and could not be evicted
`CPLB_PROT_VIOL`	A protection violation occurred

Example

```
#include <cplbtab.h>

void replace_dcplb(void) {
     int r = cplb_mgr(CPLB_EVT_DCPLB_MISS, __cplb_ctrl);
     if (r == CPLB_RELOADED)
           printf("Success\n");
     else
           printf("Failed to replace Data CPLB\n");
}
```

See also `Blackfin\lib\src\libc\crt\cplbhdr.s` in the VisualDSP++ installation directory.

See Also

cplb_hdr, cplb_init

ctime

convert calendar time into a string

Synopsis

```
#include <time.h>
char *ctime(const time_t *t);
```

Description

The `ctime` function converts a calendar time, pointed to by the argument `t` into a string that represents the local date and time. The form of the string is the same as that generated by `asctime`, and so a call to `ctime` is equivalent to

```
asctime(localtime(&t))
```

A pointer to the string is returned by `ctime`, and it may be overwritten by a subsequent call to the function.

Error Conditions

The `ctime` function does not return an error condition.

Example

```
#include <time.h>
#include <stdio.h>

time_t cal_time;

if (cal_time != (time_t)-1);
    printf("Date and Time is %s",ctime(&cal_time));
```

See Also

asctime, gmtime, localtime, time

difftime

difference between two calendar times

Synopsis

```
#include <time.h>
double difftime(time_t t1, time_t t0);
```

Description

The `difftime` function returns the difference in seconds between two calendar times, expressed as a `double`. By default, the `double` data type represents a 32-bit, single precision, floating-point, value. This form is normally insufficient to preserve all of the bits associated with the difference between two calendar times, particularly if the difference represents more than 97 days. It is recommended therefore that any function that calls `difftime` is compiled with the `-double-size-64` switch.

Error Conditions

The `difftime` function does not return an error condition.

Example

```
#include <time.h>
#include <stdio.h>
#define NA ((time_t)(-1))

time_t cal_time1;
time_t cal_time2;
double time_diff;

if ((cal_time1 == NA) || (cal_time2 == NA))
    printf("calendar time difference is not available\n");
else
    time_diff = difftime(cal_time2,cal_time1);
```

See Also

time

disable_data_cache

disable processor data caches and CPLBs

Synopsis

```
#include <cplbtab.h>
void disable_data_cache(void);
```

Description

The `disable_data_cache` function disables the processor's data caches and Data CPLBs.

The `disable_data_cache` function does not flush back to memory any modified data in the cache that is cached in Write Back mode. To flush any such data, use the `flush_data_cache` or `flush_data_buffer` routines.

Error Conditions

The `disable_data_cache` function does not return an error code.

Example

```
#include <cplbtab.h>

void cache_off(void)
{
      disable_data_cache();
}
```

See Also

cache_invalidate, enable_data_cache, flush_data_cache

div

division

Synopsis

```
#include <stdlib.h>
div_t div (int numer, int denom);
```

Description

The `div` function divides `numer` by `denom`, both of type `int`, and returns a structure of type `div_t`. The type `div_t` is defined as:

```
typedef struct {
   int quot;
   int rem;
} div_t;
```

where `quot` is the quotient of the division and `rem` is the remainder, such that if `result` is of type `div_t` , then

```
result.quot * denom + result.rem == numer
```

Error Conditions

If `denom` is zero, the behavior of the `div` function is undefined.

Example

```
#include <stdlib.h>
div_t result;

result = div(5, 2);       /* result.quot=2, result.rem=1 */
```

See Also

ldiv, fmod, modf

enable_data_cache

turn on one or both data caches

Synopsis

```
#include <cplbtab.h>
void enable_data_cache(int bitmask);
```

Description

The `enable_data_cache` function enables one or both of the processor's data caches, as indicated by the bitmask parameter:

Bit set	Meaning
`CPLB_ENABLE_CPLBS` or `CPLB_ENABLE_DCPLBS`	Enable Data CPLBs
`CPLB_ENABLE_DCACHE`	Enable data cache A
`CPLB_ENABLE_DCACHE2`	Enable data cache B

The bits are set in the following order:

1. The `CPLB_ENABLE_CPLBS` or `CPLB_ENABLE_DCPLBS` bits must be set. If neither bit is set, the function exits without changing `DMEM_CONTROL`. If one or both of these bits is set, Data CPLBs are enabled.
2. If caching is to be enabled, `CPLB_ENABLE_DCACHE` must be set. If so, Data Cache A is enabled. However, `CPLB_ENABLE_CPLBS` or `CPLB_ENABLE_DCPLBS` must be set first.
3. If both data caches are to be enabled, `CPLB_ENABLE_DCACHE2` must also be set; if so, both Data Cache A and Data Cache B are enabled. `CPLB_ENABLE_DCACHE2` is ignored if `CPLB_ENABLE_DCACHE` is not set, as the processor does not support Data Cache B in isolation.

Valid Data CPLBs must already be installed in the DCPLB active table before calling this function; the default start-up code ensures this is the case. If DCPLB Misses are a possibility, a suitable exception handler, such as `cplb_hdr`, must be installed by the default start-up code.

The data caches may only be enabled if their memory space has been left free for cache use. If any data has been mapped to this space by the `.LDF` file, the data will be corrupted by the cache's operation, and undefined behavior will result.

Error Conditions

The `enable_data_cache` function does not return an error code.

Example

```
#include <cplbtab.h>

void cache_on(int howmany)
{
      int bitmask = __cplb_ctrl;
      if (howmany == 1)
          bitmask &= ~CPLB_ENABLE_DCACHE2;
      enable_data_cache(bitmask);
}
```

See Also

cplb_hdr, cplb_init, disable_data_cache

exit

normal program termination

Synopsis

```
#include <stdlib.h>
void exit (int status);
```

Description

The `exit` function causes normal program termination. The functions registered by the `atexit` function are called in reverse order of their registration and the processor is put into the `IDLE` state. The `status` argument is stored in register `R0`, and control is passed to the label `___lib_prog_term`, which is defined by this function.

Error Conditions

The `exit` function does not return an error condition.

Example

```
#include <stdlib.h>

exit(EXIT_SUCCESS);
```

See Also

atexit, abort

exp

exponential

Synopsis

```
#include <math.h>

double exp (double x);
float expf (float x);
long double expd (long double x);
```

Description

The exp functions compute the exponential value `e` to the power of their argument.

Error Conditions

The input argument `x` for `expf` must be in the domain [-87.33, 88.72] and the input argument for `expd` must be in the domain [-708.39, 709.78]. The functions return `HUGE_VAL` if `x` is greater than the domain and 0.0 if `x` is less than the domain.

Example

```
#include <math.h>
double y;
float x;

y = exp (1.0);      /* y = 2.71828 */
x = expf (1.0);     /* x = 2.71828 */
```

See Also

log, pow

fabs

absolute value

Synopsis

```
#include <math.h>

double fabs (double x);
float fabsf (float x);
long double fabsd (long double x);
```

Description

The fabs functions return the absolute value of the argument `x`.

Error Conditions

The fabs functions do not return error conditions.

Example

```
#include <math.h>
double y;
float x;

y = fabs (-2.3);      /* y = 2.3 */
y = fabs (2.3);       /* y = 2.3 */
x = fabsf (-5.1);     /* x = 5.1 */
```

See Also

abs, labs

floor

floor

Synopsis

```
#include <math.h>

double floor (double x);
float floorf (float x);
long double floord (long double x);
```

Description

The floor functions return the largest integral value that is not greater than their argument.

Error Conditions

The floor functions do not return error conditions.

Example

```
#include <math.h>
double y;
float z;

y = floor (1.25);     /* y = 1.0  */
y = floor (-1.25);    /* y = -2.0 */
z = floorf (10.1);    /* z = 10.0 */
```

See Also

ceil

flush_data_cache

flush modified data from cache to memory

Synopsis

```
#include <cplbtab.h>
void flush_data_cache(void);
void flush_data_buffer(void *start, void *end, int invalidate);
```

Description

The `flush_data_cache` function may be used when the processor's data caches are enabled, and some data is being cached in Write Back mode. In this mode, modified data is held in the cache, and is not written back to memory immediately, thus saving the cost of an external memory access. When data is cached in this mode, it may be necessary to ensure that any modified data has been flushed to memory, so that external systems can access it. DMA transfers and dual-core accesses are common cases where Write Back mode data would need to be flushed.

The `flush_data_cache` function flushes all modified data from the cache to memory. It does so by traversing the table of active Data CPLBs, looking for valid entries that indicate a Write Back page that has been modified (that is, the Dirty flag has been set). For each page encountered, the function flushes the modified data in the page.

The `flush_data_cache` function is used by the `cplbmgr` function to flush any modified data when evicting a Dirty Write Back Data CPLB due to a Data CPLB Miss exception.

The `flush_data_buffer` function may be used when individual areas of memory need to be flushed from the cache. The function flushes data cache addresses from `start` to `end`, inclusively. Additional data addresses may also be flushed, since the function flushes entire cache lines.

If the `invalidate` parameter is non-zero, `flush_data_buffer` also invalidates the cache entries, forcing the next data access to re-fetch the data from memory. This is useful if the buffer is being updated by an external activity, such as DMA.

Error Conditions

The `flush_data_cache` function does not return an error condition.

Example

```
#include <cplbtab.h>

void do_flush(void)
{
   if (__cplb_ctrl & (CPLB_ENABLE_DCACHE|CPLB_ENABLE_DCACHE2))
         flush_data_cache();
}

char *buffer;
int buffer_len;
void inv_buffer(void) {
   flush_data_buffer(buffer, buffer+buffer_len, 1);
}
```

See Also

cache_invalidate, cplb_mgr

fmod

floating-point modulus

Synopsis

```
#include <math.h>

double fmod (double x, double y);
float fmodf (float x, float y);
long double fmodd (long double x, long double y);
```

Description

The fmod functions compute the floating-point remainder that results from dividing the first argument by the second argument.

The result is less than the second argument and has the same sign as the first argument. If the second argument is equal to zero, the fmod functions return zero.

Error Conditions

The fmod functions do not return an error condition.

Example

```
#include <math.h>
double y;
float x;

y = fmod (5.0, 2.0);      /* y = 1.0 */
x = fmodf (4.0, 2.0);     /* x = 0.0 */
```

See Also

div, ldiv, modf

free

deallocate memory

Synopsis

```
#include <stdlib.h>
void free(void *ptr);
```

Description

The `free` function deallocates a pointer previously allocated to a range of memory (by `calloc` or `malloc`) to the free memory heap. If the pointer was not previously allocated by `calloc`, `malloc` or `realloc`, the behavior is undefined.

The `free` function returns the allocated memory to the heap from which it was allocated.

Error Conditions

The `free` function does not return an error condition.

Example

```
#include <stdlib.h>
char *ptr;

ptr = (char *)malloc(10);  /* Allocate 10 bytes from heap */
free(ptr);                 /* Return space to free heap   */
```

See Also

calloc, malloc, realloc

frexp

separate fraction and exponent

Synopsis

```
#include <math.h>

double frexp(double f, int *expptr);
float frexpf (float f, int *expptr);
long double frexpd (long double f, int *expptr);
```

Description

The frexp functions separate a floating-point input into a normalized fraction and a (base 2) exponent. The functions return the first argument as a fraction which is in the interval [½, 1), and store a power of 2 in the integer pointed to by the second argument. If the input is zero, then the fraction and exponent are both set to zero.

Error Conditions

The frexp functions do not return an error condition.

Example

```
#include <math.h>
double y;
float x;
int exponent;

y = frexp (2.0, &exponent);      /* y = 0.5, exponent = 2  */
x = frexpf (4.0, &exponent);     /* x = 0.5, exponent = 3  */
```

See Also

modf

gmtime

convert calendar time into broken-down time as UTC

Synopsis

```
#include <time.h>
struct tm *gmtime(const time_t *t);
```

Description

The `gmtime` function converts a pointer to a calendar time into a broken-down time in terms of Coordinated Universal Time (UTC). A broken-down time is a structured variable, which is described in "time.h" on page 3-24.

The broken-down time is returned by `gmtime` as a pointer to static memory, which may be overwritten by a subsequent call to either `gmtime`, or to `localtime`.

Error Conditions

The `gmtime` function does not return an error condition.

Example

```
#include <time.h>
#include <stdio.h>

time_t cal_time;
struct tm *tm_ptr;

cal_time = time(NULL);
if (cal_time != (time_t) -1) {
    tm_ptr = gmtime(&cal_time);
    printf("The year is %4d\n",1900 + (tm_ptr->tm_year));
}
```

See Also

localtime, mktime, time

heap_calloc

allocate and initialize memory from a heap

Synopsis

```
#include <stdlib.h>
void *heap_calloc(int heap_index, size_t nelem, size_t size);
```

Description

The `heap_calloc` function allocates an array from the heap identified by `heap_index`. The array will contain `nelem` elements, each of size bytes; the whole array will be initialized to zero.

The function returns a pointer to the array. The return value can be safely converted to an object of any type whose size is not greater than `size*nelem` bytes. The memory allocated by `calloc` may be deallocated by either the `free` or `heap_free` functions.

Note that the `userid` of a heap is not the same as the heap's index; the index of a heap is returned by the function `heap_install` or `heap_lookup`. Refer to "Using Multiple Heaps" on page 1-245 for more information on multiple run-time heaps.

Error Conditions

The `heap_calloc` function returns a null pointer if the requested memory could not be allocated.

Example

```
#include <stdlib.h>
#include <stdio.h>

int heapid = HEAP1_USERID;
int heapindex = -1;
```

```
long *alloc_array(int nels)
{
   if (heapindex < 0) {
      heapindex = heap_lookup(heapid);
      if (heapindex == -1) {
          printf("Heap %d is not defined\n",heapid);
          exit(EXIT_FAILURE);
      }
   }
   return heap_calloc(heapindex,nels,sizeof(long));
}
```

See Also

calloc, free, heap_free, heap_init, heap_install, heap_lookup, heap_malloc, heap_realloc, heap_space_unused, malloc, realloc, space_unused

heap_free

return memory to a heap

Synopsis

```
#include <stdlib.h>
void heap_free(int heap_index, void *ptr);
```

Description

The `heap_free` function deallocates the object whose address is `ptr`, provided that `ptr` is not a null pointer. If the object was not allocated by one of heap allocation routines, or if the object has been previously freed, then the behavior of the function is undefined. If `ptr` is a null pointer, then the `heap_free` function will just return.

The function does not use the `heap_index` argument; instead it identifies the heap from which the object was allocated and returns the memory to this heap. For more information on creating multiple run-time heaps, refer to "Using Multiple Heaps" on page 1-245.

Error Conditions

The `heap_free` function does not return an error condition.

Example

```
#include <stdlib.h>

extern int userid;

int heapindex = heap_lookup(userid);
char *ptr = heap_malloc(heapindex,32 * sizeof(char));
   ...
heap_free(0,ptr);
```

See Also

calloc, free, heap_calloc, heap_init, heap_install, heap_lookup, heap_malloc, heap_realloc, heap_space_unused, malloc, realloc, space_unused

heap_init

reinitialize a heap

Synopsis

```
#include <stdlib.h>
int heap_init(int index);
```

Description

The `heap_init` function reinitializes a heap, emptying the free list, and discarding all records within the heap. Because the function discards any records within the heap, it must not be used if there are any allocations on the heap that are still active and may be used in the future.

The function returns a zero if it succeeds in re-initializing the heap specified.

Error Conditions

The `heap_init` function returns a non-zero result if it failed to re-initialize the heap.

Example

```
#include <stdlib.h>
#include <stdio.h>

int heap_index = heap_lookup(USERID_HEAP);
if (!heap_init(heap_index)) {
   printf("Heap re-initialization failed\n");
}
```

See Also

calloc, free, heap_calloc, heap_free, heap_install, heap_lookup, heap_malloc, heap_realloc, heap_space_unused, malloc, realloc, space_unused

heap_install

set up a multiple heap at run-time

Synopsis

```
#include <stdlib.h>
int heap_install(void *base, size_t length, int userid);
```

Description

The `heap_install` function initializes the heap identified by the parameter `userid`. The heap will be set up at the address specified by `base` and with a size in bytes specified by `length`. The function will return the heap index for the heap once it has been successfully initialized.

The function `heap_malloc` and the associated functions, such as `heap_calloc` and `heap_realloc`, may be used to allocate memory from the heap once the heap has been initialized. Refer to "Using Multiple Heaps" on page 1-245 for more information.

Error Conditions

The `heap_install` function returns -1 if the heap was not initialized succesfully. This may occur, for example, if there is not enough space available in the `__heaps` table, if a heap with the specified `userid` already exists, or if the new heap is too small.

Example

```
#include <stdlib.h>
#include <stdio.h>

static int heapid = 0;

int setup_heap(void *at, size_t bytes)
{
   int index;
```

```
    if ( (index = heap_install(at,bytes,++heapid)) == -1) {
         printf("Failed to initialize heap with userid %d\n",heapid);
         exit(EXIT_FAILURE);
    }
    return index;
}
```

See Also

calloc, free, heap_calloc, heap_free, heap_init, heap_lookup, heap_malloc, heap_realloc, heap_space_unused, malloc, realloc, space_unused

heap_lookup

convert a `userid` to a heap index

Synopsis

```
#include <stdlib.h>
int heap_lookup(int userid);
```

Description

The `heap_lookup` function converts a `userid` to a heap index. All heaps have a userid and a heap index associated with them. Both the `userid` and the heap index are set on heap creation. The default heap has `userid 0` and heap index `0`.

The heap index is required for the functions `heap_calloc`, `heap_malloc`, `heap_realloc`, `heap_init`, and `heap_space_unused`. For more information on creating multiple run-time heaps, refer to "Using Multiple Heaps" on page 1-245.

Error Conditions

The `heap_lookup` function returns -1 if there is no heap with the specified `userid`.

Example

```
#include <stdlib.h>
#include <stdio.h>

int heap_userid = 1;
int heap_id;

if ( (heap_id = heap_lookup(heap_userid)) == -1) {
   printf("Heap %d not setup
           - will use the default heap\n",heap_userid);
   heap_id = 0;
}
```

```
char *ptr = heap_malloc(heap_id,1024);
if (ptr == NULL) {
   printf("heap_malloc failed to allocate memory\n");
}
```

See Also

calloc, free, heap_calloc, heap_free, heap_init, heap_install, heap_malloc, heap_realloc, heap_space_unused, malloc, realloc, space_unused

heap_malloc

allocate memory from a heap

Synopsis

```
#include <stdlib.h>
void *heap_malloc(int heap_index, size_t size);
```

Description

The `heap_malloc` function allocates an object of `size` bytes, from the heap with heap index `heap_index`. It returns the address of the object if successful. The return value may be used as a pointer to an object of any type whose size in bytes is not greater than `size`.

The block of memory returned is uninitialized. The memory may be deallocated with either the `free` or `heap_free` function. For more information on creating multiple run-time heaps, refer to "Using Multiple Heaps" on page 1-245.

Error Conditions

The `heap_malloc` function returns a null pointer if it was unable to allocate the requested memory.

Example

```
#include <stdlib.h>
#include <stdio.h>

int heap_index = heap_lookup(USERID_HEAP);
long *buffer;

if (heap_index < 0) {
   printf("Heap %d is not setup\n",USERID_HEAP);
   exit(EXIT_FAILURE);
}
buffer = heap_malloc(heap_index,16 * sizeof(long));
```

```
if (buffer == NULL) {
   printf("heap_malloc failed to allocate memory\n");
}
```

See Also

calloc, free, heap_calloc, heap_free, heap_init, heap_install, heap_lookup, heap_realloc, heap_space_unused, malloc, realloc, space_unused

heap_realloc

change memory allocation from a heap

Synopsis

```
#include <stdlib.h>
void *heap_realloc(int heap_index, void *ptr, size_t size);
```

Description

The `heap_realloc` function changes the size of a previously allocated block of memory. The new size of the object in bytes is specified by the argument `size`; the new object retains the values of the old object up to its original size, but any data beyond the original size will be indeterminate. The address of the object is given by the argument `ptr`. The behavior of the function is not defined if either the object has not been allocate from a heap, or if it has already been freed.

If `ptr` is a null pointer, then `heap_realloc` behaves the same as `heap_malloc`. If `ptr` is not a null pointer, and if size is zero, then `heap_realloc` behaves the same as `heap_free`.

The argument `heap_index` is only used if `ptr` is a null pointer.

If the function successfully re-allocates the object, then it will return a pointer to the new object.

Error Conditions

If `heap_realloc` cannot reallocate the memory, it returns a null pointer and the original memory associated with `ptr` will be unchanged and will still be available.

Example

```
#include <stdlib.h>
#include <stdio.h>
```

```
int heap_index = heap_lookup(USERID_HEAP);
int *buffer;
int *temp_buffer;

if (heap_index < 0) {
   printf("Heap %d is not setup\n",USERID_HEAP);
   exit(EXIT_FAILURE);
}
buffer = heap_malloc(heap_index,32*sizeof(int));
if (buffer == NULL) {
   printf("heap_malloc failed to allocate memory\n");
}
...
temp_buffer = heap_realloc(0,buffer,64*sizeof(int));
if (temp_buffer == NULL) {
   printf(("heap_realloc failed to allocate memory\n");
} else {
   buffer = temp_buffer;
}
```

See Also

calloc, free, heap_calloc, heap_free, heap_init, heap_install, heap_lookup, heap_malloc, heap_space_unused, malloc, realloc, space_unused

heap_space_unused

space unused in specific heap

Synopsis

```
#include <stdlib.h>
int heap_space_unused(int idx);
```

Description

The `heap_space_unused` function returns the total free space in bytes for the heap with index `idx`.

Note that calling `heap_malloc(idx,heap_space_unused(idx))` does not allocate space because each allocated block uses more memory internally than the requested space. Note also that the free space in the heap may be fragmented, and thus not be available in one contiguous block.

Error Conditions

If a heap with heap index `idx` does not exist, this function returns -1.

Example

```
#include <stdlib.h>
int free_space;
free_space = heap_space_unused(1);  /* Get free space in heap 1 */
```

See Also

calloc, free, heap_calloc, heap_free, heap_init, heap_install, heap_lookup, heap_malloc, heap_realloc, malloc, realloc, space_unused

interrupt

define interrupt handling

Synopsis

```
#include <signal.h>
void (*interrupt (int sig, void(*func)(int val))) (int);
```

Description

The `interrupt` function determines how a signal received during program execution is handled. The `interrupt` function executes the function pointed to by `func` at every signal `sig`. The `signal` function executes the function only once.

The `func` argument must be one of the values listed in Table 3-16. The interrupt function causes the receipt of the signal number `sig` to be handled in one of the ways shown in Table 3-16.

Table 3-16. Interrupt Handling

Func Value	Action
`SIG_DFL`	The signal is enabled, but ignored when it occurs.
`SIG_IGN`	The signal is disabled.
`Function address`	The signal is enabled, and the function is called when the signal occurs.

The function pointed to by `func` is executed each time the interrupt is received. The `interrupt` function must be called with the `SIG_IGN` argument to disable interrupt handling. The `sig` argument may be any of the signals shown in Table 3-17 on page 3-169 which lists the supported signals in interrupt priority order from highest to lowest.

When the function pointed to by `func` is executed, the parameter `val` is set to the number of the signal that has been received. So if `func` is a signal handler used for various signals, `func` can find out which signal it is handling.

The function pointed to by `func` must not be defined using `#pragma interrupt`; the `#pragma interrupt` functions are registered using `register_handler()` instead. Refer to "Interrupt Handler Support" on page 1-201 for more information.

See Also

raise, register_handler, signal

isalnum

detect alphanumeric character

Synopsis

```
#include <ctype.h>
int isalnum(int c);
```

Description

The `isalnum` function determines whether the argument is an alphanumeric character (`A-Z`, `a-z`, or `0-9`). If the argument is not alphanumeric, `isalnum` returns a zero. If the argument is alphanumeric, `isalnum` returns a nonzero value.

Error Conditions

The `isalnum` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   printf("%3s", isalnum(ch) ? "alphanumeric" : "");
   putchar('\n');
}
```

See Also

isalpha, isdigit

isalpha

detect alphabetic character

Synopsis

```
#include <ctype.h>
int isalpha(int c);
```

Description

The `isalpha` function determines whether the input is an alphabetic character (`A-Z` or `a-z`). If the input is not alphabetic, `isalpha` returns a zero. If the input is alphabetic, `isalpha` returns a nonzero value.

Error Conditions

The `isalpha` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   printf("%2s", isalpha(ch) ? "alphabetic" : "");
   putchar('\n');
}
```

See Also

isalnum, isdigit

iscntrl

detect control character

Synopsis

```
#include <ctype.h>
int iscntrl(int c);
```

Description

The `iscntrl` function determines whether the argument is a control character (`0x00-0x1F` or `0x7F`). If the argument is not a control character, `iscntrl` returns a zero. If the argument is a control character, `iscntrl` returns a nonzero value.

Error Conditions

The `iscntrl` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   printf("%2s", iscntrl(ch) ? "control" : "");
   putchar('\n');
}
```

See Also

isalnum, isgraph

isdigit

detect decimal digit

Synopsis

```
#include <ctype.h>
int isdigit(int c);
```

Description

The `isdigit` function determines whether the input character is a decimal digit (`0-9`). If the input is not a digit, `isdigit` returns a zero. If the input is a digit, `isdigit` returns a nonzero value.

Error Conditions

The `isdigit` function does not return an error condition.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   printf("%2s", isdigit(ch) ? "digit" : "");
   putchar('\n');
}
```

See Also

isalnum, isalpha, isxdigit

isgraph

detect printable character, not including white space

Synopsis

```
#include <ctype.h>
int isgraph(int c);
```

Description

The `isgraph` function determines whether the argument is a printable character, not including white space (`0x21-0x7e`). If the argument is not a printable character, `isgraph` returns a zero. If the argument is a printable character, `isgraph` returns a nonzero value.

Error Conditions

The `isgraph` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   printf("%2s", isgraph(ch) ? "graph" : "");
   putchar('\n');
}
```

See Also

isalnum, iscntrl, isprint

isinf

test for infinity

Synopsis

```
#include <math.h>
int isinff(float x);
int isinf(double x);
int isinfd (long double x);
```

Description

The isinf functions return a zero if the argument is not set to the IEEE constant for `+Infinity` or `-Infinity`; otherwise, the functions will return a nonzero value.

Error Conditions

The isinf functions do not return or set any error conditions.

Example

```
#include <stdio.h>
#include <math.h>

static int fail=0;

main(){
   /* test int isinf(double) */
   union {
            double d; float f; unsigned long l;
   } u;

   #ifdef __DOUBLES_ARE_FLOATS__
      u.l=0xFF800000L; if ( isinf(u.d)==0 ) fail++;
      u.l=0xFF800001L; if ( isinf(u.d)!=0 ) fail++;
      u.l=0x7F800000L; if ( isinf(u.d)==0 ) fail++;
      u.l=0x7F800001L; if ( isinf(u.d)!=0 ) fail++;
```

```
    #endif

    /* test int isinff(float) */
       u.l=0xFF800000L; if ( isinff(u.f)==0 ) fail++;
       u.l=0xFF800001L; if ( isinff(u.f)!=0 ) fail++;
       u.l=0x7F800000L; if ( isinff(u.f)==0 ) fail++;
       u.l=0x7F800001L; if ( isinff(u.f)!=0 ) fail++;

    /* print pass/fail message */
    if ( fail==0 )
       printf("Test passed\n");
    else
       printf("Test failed: %d\n", fail);
}
```

See Also

isnan

islower

detect lowercase character

Synopsis

```
#include <ctype.h>
int islower(int c);
```

Description

The `islower` function determines whether the argument is a lowercase character (`a-z`). If the argument is not lowercase, `islower` returns a zero. If the argument is lowercase, `islower` returns a nonzero value.

Error Conditions

The `islower` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   printf("%2s", islower(ch) ? "lowercase" : "");
   putchar('\n');
}
```

See Also

isalpha, isupper

isnan

test for Not a Number (NAN)

Synopsis

```
#include <math.h>
int isnanf(float x);
int isnan(double x);
int isnand (long double x);
```

Description

The isnan functions return a zero if the argument is not set to an IEEE `NaN`; otherwise, the functions return a nonzero value.

Error Conditions

The isnan functions do not return or set any error conditions.

Example

```
#include <stdio.h>
#include <math.h>

static int fail=0;

main(){
   /* test int isnan(double) */
   union {
      double d; float f; unsigned long l;
   } u;

   #ifdef __DOUBLES_ARE_FLOATS__
      u.l=0xFF800000L; if ( isnan(u.d)!=0 ) fail++;
      u.l=0xFF800001L; if ( isnan(u.d)==0 ) fail++;
      u.l=0x7F800000L; if ( isnan(u.d)!=0 ) fail++;
      u.l=0x7F800001L; if ( isnan(u.d)==0 ) fail++;
   #endif
```

```
    /* test int isnanf(float) */
       u.l=0xFF800000L; if ( isnanf(u.f)!=0 ) fail++;
       u.l=0xFF800001L; if ( isnanf(u.f)==0 ) fail++;
       u.l=0x7F800000L; if ( isnanf(u.f)!=0 ) fail++;
       u.l=0x7F800001L; if ( isnanf(u.f)==0 ) fail++;

    /* print pass/fail message */
    if ( fail==0 )
       printf("Test passed\n");
    else
       printf("Test failed: %d\n", fail);
}
```

See Also

isinf

isprint

detect printable character

Synopsis

```
#include <ctype.h>
int isprint(int c);
```

Description

The `isprint` function determines whether the argument is a printable character (`0x20-0x7E`). If the argument is not a printable character, `isprint` returns a zero. If the argument is a printable character, `isprint` returns a nonzero value.

Error Conditions

The `isprint` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   printf("%3s", isprint(ch) ? "printable" : "");
   putchar('\n');
}
```

See Also

isgraph, isspace

ispunct

detect punctuation character

Synopsis

```
#include <ctype.h>
int ispunct(int c);
```

Description

The `ispunct` function determines whether the argument is a punctuation character. If the argument is not a punctuation character, `ispunct` returns a zero. If the argument is a punctuation character, `ispunct` returns a non-zero value.

Error Conditions

The `ispunct` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   printf("%3s", ispunct(ch) ? "punctuation" : "");
   putchar('\n');
}
```

See Also

isalnum

isspace

detect whitespace character

Synopsis

```
#include <ctype.h>
int isspace(int c);
```

Description

The `isspace` function determines whether the argument is a blank whitespace character (`0x09-0x0D` or `0x20`). This includes the characters space ( ), form feed (`\f`), new line (`\n`), carriage return (`\r`), horizontal tab (`\t`), and vertical tab (`\v`).

If the argument is not a blank space character, `isspace` returns a zero. If the argument is a blank space character, `isspace` returns a nonzero value.

Error Conditions

The `isspace` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   printf("%2s", isspace(ch) ? "space" : "");
   putchar('\n');
}
```

See Also

iscntrl, isgraph

isupper

detect uppercase character

Synopsis

```
#include <ctype.h>
int isupper(int c);
```

Description

The `isupper` function determines whether the argument is an uppercase character (`A-Z`). If the argument is not an uppercase character, `isupper` returns a zero. If the argument is an uppercase character, `isupper` returns a nonzero value.

Error Conditions

The `isupper` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   printf("%2s", isupper(ch) ? "uppercase" : "");
   putchar('\n');
}
```

See Also

isalpha, islower

isxdigit

detect hexadecimal digit

Synopsis

```
#include <ctype.h>
int isxdigit(int c);
```

Description

The `isxdigit` function determines whether the argument is a hexadecimal digit character (`A-F`, `a-f`, or `0-9`). If the argument is not a hexadecimal digit, `isxdigit` returns a zero. If the argument is a hexadecimal digit, `isxdigit` returns a nonzero value.

Error Conditions

The `isxdigit` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   printf("%2s", isxdigit(ch) ? "hexadecimal" : "");
   putchar('\n');
}
```

See Also

isalnum, isdigit

l1_memcpy, memcpy_l1

Copy instructions from L1 Instruction Memory to data memory, and from data memory to L1 Instruction Memory.

Synopsis

```
#include <ccblkfn.h>
void *_l1_memcpy(void *datap, const void *instrp, size_t n);
void *_memcpy_l1(void *instrp, const void *datap, size_t n);
```

Description

The `_l1_memcpy` function copies `n` characters of program instructions from the address `instrp` to the data buffer `datap`. The `_memcpy_l1` function is the inverse: it copies `n` characters of program instructions from the data buffer `datap` to the address `instrp`. For both functions, there are a number of restrictions:

- `n` must be a multiple of 8
- `instrp` must be an address in L1 Instruction Memory
- `instrp` must be 8-byte aligned
- `datap` must be 4-byte aligned
- `instrp+n-1` must be within L1 Instruction Memory
- for dual-core processors, `instrp` must correspond to the core calling the function.

The `_l1_memcpy` function returns `datap` for success. The `_memcpy_l1` function returns `instrp` for success.

The C++ run-time library uses `_l1_memcpy` to implement the C++ exception throw/catch mechanism. The C and C++ run-time libraries use `_memcpy_l1` to implement the memory-initialization process, if the `.DXE` file has been built with the `-mem` compiler switch, or with the `-meminit` linker switch.

Error Conditions

If any of the restrictions are not met, the `_l1_memcpy` and `_memcpy_l1` functions return NULL.

Examples

```
/* copying program instructions from L1 Instruction
** memory to data memory.
*/
#include <ccblkfn.h>
char dest[32];
const char *src = (const char *)0xFFA00000;
if (_l1_memcpy(dest, src, 32) != dest)
   exit(1);

/* copying program instructions from data memory
** to L1 Instruction memory.
*/
#include <ccblkfn.h>
const char src[32] = { /* some instruction op-codes */ };
char *dest = (char *)0xFFA00000;
if (_memcpy_l1(dest, src, 32) != dest)
   exit(1);
```

See Also

memcpy

labs

long integer absolute value

Synopsis

```
#include <stdlib.h>
long int labs(long int j);
long long int llabs (long long int j);
```

Description

The `labs` and `llabs` functions return the absolute value of their integer inputs.

Error Conditions

The `labs` and `llabs` functions do not return an error condition.

Example

```
#include <stdlib.h>
long int j;
j = labs(-285128);      /* j = 285128 */
```

See Also

abs, fabs

ldexp

multiply by power of 2

Synopsis

```
#include <math.h>
double ldexp (double x, int n);
float ldexpf (float x, int n);
long double ldexpd (long double x, int n);
```

Description

The ldexp functions return the value of the floating-point argument multiplied by 2^n. These functions add the value of `n` to the exponent of `x`.

Error Conditions

If the result overflows, the ldexp functions return `HUGE_VAL` with the proper sign and set `errno` to `ERANGE`. If the result underflows, a zero is returned.

Example

```
#include <math.h>
double y;
float x;

y = ldexp (0.5, 2);       /* y = 2.0  */
x = ldexpf (1.0, 2);      /* x = 4.0  */
```

See Also

exp, pow

ldiv

long division

Synopsis

```
#include <stdlib.h>
ldiv_t ldiv(long int numer, long int denom);
lldiv_t lldiv (long long int numer, long long int denom);
```

Description

The `ldiv` and `lldiv` functions divide `numer` by `denom` and return a structure of type `ldiv_t` and `lldiv_t`, respectively. The types `ldiv_t` and `lldiv_t` are defined as:

```
typedef struct {
   long int quot;
   long int rem;
} ldiv_t;

typedef struct {
   long long int quot;
   long long int rem;
} lldiv_t;
```

where `quot` is the quotient of the division and `rem` is the remainder, such that if `result` is of the appropriate type, then

```
result.quot * denom + result.rem = numer
```

Error Conditions

If `denom` is zero, the behavior of the `ldiv` and `lldiv` functions are undefined.

Example

```
#include <stdlib.h>
ldiv_t result;

result = ldiv(7, 2);      /* result.quot=3, result.rem=1 */
```

See Also

div, fmod

localtime

convert calendar time into broken-down time

Synopsis

```
#include <time.h>
struct tm *localtime(const time_t *t);
```

Description

The `localtime` function converts a pointer to a calendar time into a broken-down time that corresponds to current time zone. A broken-down time is a structured variable, which is described in "time.h" on page 3-24. This implementation of the header file does not support the Daylight Saving flag nor does it support time zones and, thus, `localtime` is equivalent to the `gmtime` function.

The broken-down time is returned by `localtime` as a pointer to static memory, which may be overwritten by a subsequent call to either `localtime`, or to `gmtime`.

Error Conditions

The `localtime` function does not return an error condition.

Example

```
#include <time.h>
#include <stdio.h>

time_t cal_time;
struct tm *tm_ptr;

cal_time = time(NULL);
if (cal_time != (time_t) -1) {
    tm_ptr = localtime(&cal_time);
    printf("The year is %4d\n",1900 + (tm_ptr->tm_year));
}
```

See Also

asctime, gmtime, mktime, time

log

natural logarithm

Synopsis

```
#include <math.h>
double log (double x);
float logf (float x);
long double logd (long double x);
```

Description

The log functions compute the natural (base e) logarithm of their argument.

Error Conditions

The log functions return `-HUGE_VAL` if the input value is zero or negative.

Example

```
#include <math.h>
double y;
float x;

y = log (1.0);          /* y = 0.0 */
x = logf (2.71828);     /* x = 1.0 */
```

See Also

alog, exp, log10

log10

base 10 logarithm

Synopsis

```
#include <math.h>
double log10(double f);
float log10f (float f);
long double log10d (long double f);
```

Description

The log10 functions return the base 10 logarithm of their inputs.

Error Conditions

The log10 functions return `-HUGE_VAL` if the input is zero or negative.

Example

```
#include <math.h>
double y;
float x;

y = log10 (100.0);     /* y = 2.0 */
x = log10f (10.0);     /* x = 1.0 */
```

See Also

alog10, log, pow

longjmp

second return from `setjmp`

Synopsis

```
#include <setjmp.h>
void longjmp(jmp_buf env, int return_val);
```

Description

The `longjmp` function causes the program to execute a second return from the place where `setjmp (env)` was called (with the same `jmp_buf` argument).

The `longjmp` function takes as its arguments a jump buffer that contains the context at the time of the original call to `setjmp`. It also takes an integer, `return_val`, which `setjmp` returns if `return_val` is nonzero. Otherwise, `setjmp` returns a 1.

If `env` was not initialized through a previous call to `setjmp` or the function that called `setjmp` has since returned, the behavior is undefined. Also, automatic variables that are local to the original function calling `setjmp`, that do not have `volatile` qualified type and that have changed their value prior to the `longjmp` call, have indeterminate value.

Error Conditions

The `longjmp` function does not return an error condition.

Example

```
#include <setjmp.h>
#include <stdio.h>
#include <errno.h>
#include <stdlib.h>
```

```
jmp_buf env;
int res;

if ((res == setjmp(env)) != 0) {
   printf ("Problem %d reported by func ()", res);
   exit (EXIT_FAILURE);
}
func ();

void func (void)
{
   if (errno != 0) {
      longjmp (env, errno);
   }
}
```

See Also

setjmp

malloc

allocate memory

Synopsis

```
#include <stdlib.h>
void *malloc(size_t size);
```

Description

The `malloc` function returns a pointer to a block of memory of length `size`. The block of memory is not initialized.

Error Conditions

The `malloc` function returns a null pointer if it is unable to allocate the requested memory.

Example

```
#include <stdlib.h>
long *ptr;

ptr = (long *)malloc(10 * sizeof(long));    /* ptr points to an */
                                            /* array of 10 longs */
```

See Also

calloc, realloc, free

memchr

find first occurrence of character

Synopsis

```
#include <string.h>
void *memchr(const void *s1, int c, size_t n);
```

Description

The `memchr` function compares the range of memory pointed to by `s1` with the input character `c`, and returns a pointer to the first occurrence of `c`. A null pointer is returned if `c` does not occur in the first `n` characters.

Error Conditions

The `memchr` function does not return an error condition.

Example

```
#include <string.h>
char *ptr;

ptr= memchr("TESTING", 'E', 7);
     /* ptr points to the E in TESTING */
```

See Also

strchr, strrchr

memcmp

compare objects

Synopsis

```
#include <string.h>
int memcmp(const void *s1, const void *s2, size_t n);
```

Description

The `memcmp` function compares the first `n` characters of the objects pointed to by `s1` and `s2`. This function returns a positive value if the `s1` object is lexically greater than the `s2` object, a negative value if the `s2` object is lexically greater than the `s1` object, and a zero if the objects are the same.

Error Conditions

The `memcmp` function does not return an error condition.

Example

```
#include <string.h>
char string1 = "ABC";
char string2 = "BCD";
int result;

result = memcmp (string1, string2, 3);      /* result < 0 */
```

See Also

strcmp, strcoll, strncmp

memcpy

copy characters from one object to another

Synopsis

```
#include <string.h>
void *memcpy(void *s1, const void *s2, size_t n);
```

Description

The `memcpy` function copies `n` characters from the object pointed to by `s2` into the object pointed to by `s1`. The behavior of `memcpy` is undefined if the two objects overlap.

The `memcpy` function returns the address of `s1`.

Error Conditions

The `memcpy` function does not return an error condition.

Example

```
#include <string.h>
char *a = “SRC”;
char *b = “DEST”;
memcpy (b, a, 3);       /*  b=”SRCT”  */
```

See Also

memmove, strcpy, strncpy

memmove

copy characters between overlapping objects

Synopsis

```
#include <string.h>
void *memmove(void *s1, const void *s2, size_t n);
```

Description

The `memmove` function copies `n` characters from the object pointed to by `s2` into the object pointed to by `s1`. The entire object is copied correctly even if the objects overlap.

The `memmove` function returns a pointer to `s1`.

Error Conditions

The `memmove` function does not return an error condition.

Example

```
#include <string.h>
char *ptr, *str = "ABCDE";

ptr = str + 2;
memmove(ptr, str, 3);      /*  ptr = "ABC", str = "ABABC"  */
```

See Also

memcpy, strcpy, strncpy

memset

set range of memory to a character

Synopsis

```
#include <string.h>
void *memset(void *s1, int c, size_t n);
```

Description

The `memset` function sets a range of memory to the input character `c`. The first `n` characters of `s1` are set to `c`.

The `memset` function returns a pointer to `s1`.

Error Conditions

The `memset` function does not return an error condition.

Example

```
#include <string.h>
char string1[50];
memset(string1, '\0', 50);      /* set string1 to 0 */
```

See Also

memcpy

mktime

convert broken-down time into a calendar time

Synopsis

```
#include <time.h>
time_t mktime(struct tm *tm_ptr);
```

Description

The `mktime` function converts a pointer to a broken-down time, which represents a local date and time, into a calendar time. However, this implementation of `time.h` does not support either daylight saving or time zones and hence this function will interpret the argument as Greenwich Mean Time (UTC).

A broken-down time is a structured variable which is defined in the `time.h` header file as:

```
struct tm { int tm_sec;   /* seconds after the minute [0,61] */
   int tm_min;            /* minutes after the hour [0,59]   */
   int tm_hour;           /* hours after midnight [0,23]     */
   int tm_mday;           /* day of the month [1,31]         */
   int tm_mon;            /* months since January [0,11]     */
   int tm_year;           /* years since 1900                */
   int tm_wday;           /* days since Sunday [0, 6]        */
   int tm_yday;           /* days since January 1st [0,365]  */
   int tm_isdst;          /* Daylight Saving flag            */
};
```

The various components of the broken-down time are not restricted to the ranges indicated above. The `mktime` function calculates the calendar time from the specified values of the components (ignoring the initial values of `tm_wday` and `tm_yday`), and then"'normalizes" the broken-down time forcing each component into its defined range.

Error Conditions

The `mktime` function returns the value `(time_t) -1` if the calendar time cannot be represented.

Example

```
#include <time.h>
#include <stdio.h>

static const char *wday[] = {"Sun","Mon","Tue","Wed",
                             "Thu","Fri","Sat","???"};

struct tm tm_time = {0,0,0,0,0,0,0,0,0};

tm_time.tm_year = 2000 - 1900;
tm_time.tm_mday = 1;

if (mktime(&tm_time) == -1)
    tm_time.tm_wday = 7;
printf("%4d started on a %s\n",
    1900 + tm_time.tm_year,
    wday[tm_time.tm_wday]);
```

See Also

gmtime, localtime, time

modf

separate integral and fractional parts

Synopsis

```
#include <math.h>
double modf (double x, double *intptr);
float modff (float x, float *intptr);
long double modfd (long double x, long double *intptr);
```

Description

The modf functions separate the first argument into integral and fractional portions. The fractional portion is returned and the integral portion is stored in the object pointed to by `intptr`. The integral and fractional portions have the same sign as the input.

Error Conditions

The modf functions do not return error conditions.

Example

```
#include <math.h>
double y, n;
float m, p;

y = modf (-12.345, &n);     /* y = -0.345, n = -12.0 */
m = modff (11.75, &p);      /* m = 0.75, p = 11.0    */
```

See Also

frexp

pow

raise to a power

Synopsis

```
#include <math.h>
double pow (double x, double y);
float powf (float x, float y);
long double powd (long double x, long double y);
```

Description

The pow functions compute the value of the first argument raised to the power of the second argument.

Error Conditions

The functions return zero when the first argument `x` is zero and the second argument `y` is not an integral value. When `x` is zero and `y` is less than zero, or when the result cannot be represented, the functions will return the constant `HUGE_VAL`.

Example

```
#include <math.h>
double z;
float x;

z = pow (4.0, 2.0);      /* z = 16.0 */
x = powf (4.0, 2.0);     /* x = 16.0 */
```

See Also

exp, ldexp

qsort

quicksort

Synopsis

```
#include <stdlib.h>
void qsort (void *base, size_t nelem, size_t size,
            int (*compare) (const void *, const void *));
```

Description

The `qsort` function sorts an array of `nelem` objects, pointed to by `base`. Each object is specified by its `size`.

The contents of the array are sorted into ascending order according to a comparison function pointed to by `compare`, which is called with two arguments that point to the objects being compared. The function returns an integer less than, equal to, or greater than zero if the first argument is considered to be respectively less than, equal to, or greater than the second.

If two elements compare as equal, their order in the sorted array is unspecified. The `qsort` function executes a binary search operation on a pre-sorted array. Note that:

- `base` points to the start of the array
- `nelem` is the number of elements in the array
- `size` is the size of each element of the array
- `compare` is a pointer to a function that is called by `qsort` to compare two elements of the array. The function returns a value less than, equal to, or greater than zero, according to whether the first argument is less than, equal to, or greater than the second.

Error Condition

The `qsort` function returns no value.

Example

```
#include <stdlib.h>
float a[10];

int compare_float (const void *a, const void *b)
{
   float aval = *(float *)a;
   float bval = *(float *)b;
   if (aval < bval)
      return -1;
   else if (aval == bval)
      return 0;
   else
      return 1;
}
qsort (a, sizeof (a)/sizeof (a[0]), sizeof (a[0]),compare_float);
```

See Also

bsearch

raise

force a signal

Synopsis

```
#include <signal.h>
int raise(int sig);
```

Description

The `raise()` function sends the signal `sig` to the executing program. The `raise` function forces interrupts wherever possible and simulates an interrupt otherwise. The `sig` argument must be one of the signals listed in priority order in Table 3-17.

Table 3-17. Raise Function Signals – Values and Meanings

Sig Value	Definition
`SIGEMU`	emulation trap
`SIGRSET`	machine reset
`SIGNMI`	non-maskable interrupt
`SIGEVNT`	event vectoring
`SIGHW`	hardware error
`SIGTIMR`	timer events Note that `SIGALRM` is mapped onto the signal `SIGTIMR`
`SIGIVG7 - SIGIVG15`	miscellaneous interrupts Note that:`SIGUSR1` is mapped onto the signal `SIGIVG15` `SIGUSR2` is mapped onto the signal `SIGIVG14`
`SIGINT`	software interrupt
`SIGILL`	software interrupt
`SIGBUS`	software interrupt
`SIGFPE`	software interrupt

Table 3-17. Raise Function Signals – Values and Meanings (Cont'd)

Sig Value	Definition
SIGSEGV	software interrupt
SIGTERM	software interrupt
SIGABRT	software interrupt

When an interrupt is forced, the current ISR registered in the Event Vector Table is invoked. Normally, this is a dispatcher installed by `signal()`, which saves the context before invoking the signal handler, and restores it afterwards.

When an interrupt is simulated, `raise()` calls the registered signal handler directly.

Error Conditions

The `raise` function returns a zero if successful, a nonzero value if it fails.

Example

```
#include <signal.h>
raise(SIGABRT);
```

See Also

interrupt, signal

rand

random number generator

Synopsis

```
#include <stdlib.h>
int rand(void);
```

Description

The `rand` function returns a pseudo-random integer value in the range [0, 2^{30}– 1].

For this function, the measure of randomness is its periodicity—the number of values it is likely to generate before repeating a pattern. The output of the pseudo-random number generator has a period in the order of 2^{30}– 1.

Error Conditions

The `rand` function does not return an error condition.

Example

```
#include <stdlib.h>
int i;

i = rand();
```

See Also

srand

realloc

change memory allocation

Synopsis

```
#include <stdlib.h>
void *realloc(void *ptr, size_t size);
```

Description

The `realloc` function changes the memory allocation of the object pointed to by `ptr` to `size`. Initial values for the new object are taken from the values in the object pointed to by `ptr`. If the size of the new object is greater than the size of the object pointed to by `ptr`, then the values in the newly allocated section are undefined.

If `ptr` is a non-null pointer that was not allocated with `malloc` or `calloc`, the behavior is undefined. If `ptr` is a null pointer, `realloc` imitates `malloc`. If `size` is zero and `ptr` is not a null pointer, `realloc` imitates `free`.

Error Conditions

If memory cannot be allocated, `ptr` remains unchanged and `realloc` returns a null pointer.

Example

```
#include <stdlib.h>
int *ptr;

ptr = malloc(10 * sizeof(int));          /* ptr points to an array
                                            of 10 ints          */
ptr = realloc(ptr,20 * sizeof(int));     /* ptr now points to an
                                            array of 20 ints    */
```

See Also

calloc, free, malloc

register_handler

register event handlers directly

Synopsis

```
#include <sys/exception.h>
ex_handler_fn register_handler (interrupt_kind kind, ex_handler_fn fn);
```

Description

The `register_handler` function determines how the hardware event `kind` is handled. This is done by registering the function pointed to by `fn` as a handler for the event. The `kind` event is an enumeration identifying each of the hardware events—interrupts and exceptions—accepted by the Blackfin processor.

For the values for `kind`, refer to "Registering an ISR" on page 1-204. The `fn` must be one of the values listed here.

fn Value	Action
`EX_INT_IGNORE`	The event is disabled; the vector table is unchanged.
`EX_INT_DEFAULT`	The event is disabled; the vector table is cleared.
`Function address`	The event is enabled; the address is entered into the vector table.

The vector table is used by the Blackfin processor to identify instructions to execute when an event occurs. When a given event is raised, and if the event is enabled, the processor begins executing instructions from the address given by the event's entry in the vector table.

No dispatcher is used to invoke `fn`. Therefore, `fn` must be a full event handler. That is, it must save the processor context on entry, restore the context on exit, and return using the machine instruction appropriate to the event type. Therefore, if `fn` is written in C, it must be defined with an appropriate `#pragma` to ensure the compiler generates suitable code. A nor-

mal C function is *not* suitable for use with `register_handler`. The header file `<sys/exception.h>` provides macros to be used with `register_handler()` for prototyping and declaring functions.

The `register_handler` function is a more direct mechanism than `signal()` and `interrupt()`. The `signal()` and `interrupt()` functions accept (and require) "normal" C functions, and therefore need to use a dispatcher to invoke the registered function. In contrast, `register_handler()` does not use a dispatcher, and so, "normal" C functions are *not* suitable for registering with the `register_handler` function.

Note that `register_handler()` does not modify the interrupt latch register. Therefore, if `register_handler()` is called to install a handler for a latched interrupt, the interrupt handler is called during the execution of `register_handler()`. The appropriate bit in the interrupt latch register must be unset by the user if this is undesirable behavior. See the appropriate hardware reference manual for details of how to do this.

Refer to "Interrupt Handler Support" on page 1-201 for more information.

Example

```
#include <sys/exception.h>
int timer_count = 0;

EX_INTERRUPT_HANDLER(inccount)
{
    timer_count++;
}

main(void)
{
    register_handler(ik_timer, inccount);   /* . . . .  */
}
```

See Also

raise, interrupt, signal

setjmp

define a run-time label

Synopsis

```
#include <setjmp.h>
int setjmp(jmp_buf env);
```

Description

The `setjmp` function saves the calling environment in the `jmp_buf` argument. The effect of the call is to declare a run-time label that can be jumped to via a subsequent call to `longjmp`.

When `setjmp` is called, it immediately returns with a result of zero to indicate that the environment has been saved in the `jmp_buf` argument. If, at some later point, `longjmp` is called with the same `jmp_buf` argument, `longjmp` restores the environment from the argument. The execution then resumes at the statement immediately following the corresponding call to `setjmp`. The effect is as if the call to `setjmp` has returned for a second time but this time the function returns a nonzero result.

The effect of calling `longjmp` is undefined if the function that called `setjmp` has returned in the interim.

Error Conditions

The label `setjmp` does not return an error condition.

Example

See code example for "longjmp" on page 3-155.

See Also

longjmp

signal

define signal handling

Synopsis

```
#include <signal.h>
void (*signal(int sig, void (*func)(int val))) (int);
```

Description

The `signal` function determines how a signal received during program execution is handled. This function causes a response to any single occurrence of an interrupt. The `sig` argument must be one of the signals listed in priority order in Table 3-17 on page 3-169.

Event handlers may also be installed directly; for more information, refer to "Interrupt Handler Support" on page 1-201. The default run-time header installs event handlers that invoke handlers registered by `signal()`.

The `signal` function installs a dispatcher ISR into the Event Vector Table, and enables the relevant event. When the event occurs, the dispatcher saves the processor context before the invoked `func`, and restores the context afterwards.

- If the function is `SIG_DFL`, the signal is enabled, but ignored when it occurs.
- If the function is `SIG_IGN`, the signal is disabled.

When the function pointed to by `func` is executed, the parameter `val` is set to the number of the signal that has been received. Thus, it has the same value as `sig`, assuming that for each signal `sig`, a unique function is registered.

The function pointed to by `func` must not be defined using `#pragma interrupt`; the `#pragma interrupt` functions are registered using `register_handler()` instead. Refer to "Registering an ISR" on page 1-204 and "ISRs and ANSI C Signals" on page 1-204 for more information.

See Also

interrupt, raise, register_handler

sin

sine

Synopsis

```
#include <math.h>

float sinf (float x);
double sin (double x);
long double sind (long double x);

fract16 sin_fr16 (float x);
```

Description

The sine functions return the sine of the argument. Both the argument `x` and the results returned by the functions are in radians.

The result returned by the functions `sine`, `sinf`, and `sind` is in the range [-1, 1].

The `sin_fr16` function inputs a fractional value in the range [-1.0, 1.0) corresponding to [-π/2, π/2]. The domain represents half a cycle which can be used to derive a full cycle if required (see Notes below). The result, in radians, is in the range [-1.0, 1.0).

Error Conditions

The sin functions do not return an error condition.

Example

```
#include <math.h>
double y;
y = sin(4.04159);      /* y = 0.0 */
```

Notes

The domain of the `sin_fr16` function is restricted to the fractional range [`0x8000`, `0x7fff`] which corresponds to half a period from –(π /2) to π/2. It is possible however to derive the full period using the following properties of the function.

```
sine [0, π/2] = -sine [π, 3/2 π]
sine [-π/2, 0] = -sine [π/2, π]
```

The function below uses these properties to calculate the full period (from 0 to 2π) of the `sine` function using an input domain of `[0, 0x7fff]`.

```
#include <math.h>

fract16 sin2pi_fr16 (fract16 x)
{
   if (x < 0x2000) {                    /* <0.25  */
      /* first quadrant [0..π/2):                */
      /* sin_fr16([0x0..0x7fff]) = [0..0x7fff)   */
      return sin_fr16(x * 4);

   } else if (x < 0x6000) {            /* < 0.75  */
      /* if (x < 0x4000)                         */
      /* second quadrant [π/2..π):               */
      /* -sin_fr16([0x8000..0x0)) = [0x7fff..0)  */
      /*                                         */
      /* if (x < 0x6000)                         */
      /* third quadrant [π..3/2π):               */
      /* -sin_fr16([0x0..0x7fff]) = [0..0x8000)  */
      return -sin_fr16((0xc000 + x) * 4);

   } else {
      /* fourth quadrant [3/2π..π):              */
      /* sin_fr16([0x8000..0x0)) = [0x8000..0)   */
      return sin_fr16((0x8000 + x) * 4);
   }
}
```

See Also

asin, cos

sinh

hyperbolic sine

Synopsis

```
#include <math.h>

double sinh (double x);
float sinhf (float x);
long double sinhd (long double x);
```

Description

The sinh functions return the hyperbolic sine of `x`.

Error Conditions

The input argument `x` must be in the domain [-89.41, 89.41] for `sinhf`, and in the domain [-710.46, 710.47] for `sinhd`. If the input value is greater than the function's domain, then `HUGE_VAL` is returned, and if the input value is less than the domain, then `-HUGE_VAL` is returned.

Example

```
#include <math.h>
double x, y;
float z, w;

y = sinh (x);
z = sinhf (w);
```

See Also

cosh

space_unused

space unused in heap

Synopsis

```
#include <stdlib.h>
int space_unused(void);
```

Description

The `space_unused` function returns the total free space in bytes for the default heap. Note that calling `malloc(space_unused())` does not allocate space because each allocated block uses more memory internally than the requested space, and also the free space in the heap may be fragmented, and thus not be available in one contiguous block.

Error Conditions

If there are no heaps, calling this function will return -1.

Example

```
#include <stdlib.h>
int free_space;
free_space = space_unused();   /* Get free space in the heap */
```

See Also

calloc, free, heap_calloc, heap_free, heap_init, heap_install, heap_lookup, heap_malloc, heap_space_unused, malloc, realloc

sqrt

square root

Synopsis

```
#include <math.h>

float sqrtf (float x);
double sqrt (double x);
long double sqrtd (long double x);

fract16 sqrt_fr16 (fract16 x);
```

Description

The sqrt functions return the positive square root of the argument `x`.

Error Conditions

The sqrt functions return a zero if the input argument is negative.

Example

```
#include <math.h>
double y;
y = sqrt(2.0);      /* y = 1.414..... */
```

See Also

rsqrt

srand

random number seed

Synopsis

```
#include <stdlib.h>
void srand(unsigned int seed);
```

Description

The `srand` function sets the seed value for the `rand` function. A particular seed value always produces the same sequence of pseudo-random numbers.

Error Conditions

The `srand` function does not return an error condition.

Example

```
#include <stdlib.h>
srand(22);
```

See Also

rand

strcat

concatenate strings

Synopsis

```
#include <string.h>
char *strcat(char *s1, const char *s2);
```

Description

The `strcat` function appends a copy of the null-terminated string pointed to by `s2` to the end of the null-terminated string pointed to by `s1`. The function returns a pointer to the new `s1` string, which is null-terminated. The behavior of `strcat` is undefined if the two strings overlap.

Error Conditions

The `strcat` function does not return an error condition.

Example

```
#include <string.h>
char string1[50];

string1[0] = 'A';
string1[1] = 'B';
string1[2] = '\0';
strcat(string1, "CD");      /* new string is "ABCD" */
```

See Also

strncat

strchr

find first occurrence of character in string

Synopsis

```
#include <string.h>
char *strchr(const char *s1, int c);
```

Description

The `strchr` function returns a pointer to the first location in `s1` (null-terminated string) that contains the character `c`.

Error Conditions

The `strchr` function returns a null pointer if `c` is not part of the string.

Example

```
#include <string.h>
char *ptr1, *ptr2;

ptr1 = "TESTING";
ptr2 = strchr(ptr1, 'E');
       /* ptr2 points to the E in TESTING */
```

See Also

memchr, strrchr

strcmp

compare strings

Synopsis

```
#include <string.h>
int strcmp(const char *s1, const char *s2);
```

Description

The `strcmp` function lexicographically compares the null-terminated strings pointed to by `s1` and `s2`. The function returns a positive value if the `s1` string is greater than the `s2` string, a negative value if the `s2` string is greater than the `s1` string, and a zero if the strings are the same.

Error Conditions

The `strcmp` function does not return an error condition.

Example

```
#include <string.h>
char string1[50], string2[50];

if (strcmp(string1, string2))
    printf("%s is different than %s \n", string1, string2);
```

See Also

memcmp, strncmp

strcoll

compare strings

Synopsis

```
#include <string.h>
int strcoll(const char *s1, const char *s2);
```

Description

The `strcoll` function compares the string pointed to by `s1` with the string pointed to by `s2`. The comparison is based on the `LC_COLLATE` locale macro. Because only the C locale is defined in the Blackfin run-time environment, the `strcoll` function is identical to the `strcmp` function. The function returns a positive value if the `s1` string is greater than the `s2` string, a negative value if the `s2` string is greater than the `s1` string, and a zero if the strings are the same.

Error Conditions

The `strcoll` function does not return an error condition.

Example

```
#include <string.h>
char string1[50], string2[50];

if (strcoll(string1, string2))
   printf("%s is different than %s \n", string1, string2);
```

See Also

strcmp, strncmp

strcpy

copy from one string to another

Synopsis

```
#include <string.h>
void *strcpy(char *s1, const char *s2);
```

Description

The `strcpy` function copies the null-terminated string pointed to by `s2` into the space pointed to by `s1`. Memory allocated for `s1` must be large enough to hold `s2`, plus one space for the null character (‘`\0`’). The behavior of `strcpy` is undefined if the two objects overlap, or if `s1` is not large enough. The `strcpy` function returns the new `s1`.

Error Conditions

The `strcpy` function does not return an error condition.

Example

```
#include <string.h>
char string1[50];

strcpy(string1, “SOMEFUN”);
   /* SOMEFUN is copied into string1 */
```

See Also

memcpy, memmove, strncpy

strcspn

length of character segment in one string but not the other

Synopsis

```
#include <string.h>
size_t strcspn(const char *s1, const char *s2);
```

Description

The `strcspn` function returns the length of the initial segment of `s1` which consists entirely of characters not in the string pointed to by `s2`. The string pointed to by `s2` is treated as a set of characters. The order of the characters in the string is not significant.

Error Conditions

The `strcspn` function does not return an error condition.

Example

```
#include <string.h>
char *ptr1, *ptr2;
size_t len;

ptr1 = "Tried and Tested";
ptr2 = "aeiou";
len = strcspn (ptr1,ptr2);      /* len = 2 */
```

See Also

strlen, strspn

strerror

get string containing error message

Synopsis

```
#include <string.h>
char *strerror(int errnum);
```

Description

The `strerror` function returns a pointer to a string containing an error message by mapping the number in `errnum` to that string.

Error Conditions

The `strerror` function does not return an error condition.

Example

```
#include <string.h>
char *ptr1;

ptr1 = strerror(1);
```

See Also

No references to this function.

strftime

format a broken-down time

Synopsis

```
#include <time.h>
size_t strftime(char *buf,
                size_t buf_size,
                const char *format,
                const struct tm *tm_ptr);
```

Description

The `strftime` function formats the broken-down time `tm_ptr` into the `char` array pointed to by `buf`, under the control of the format string `format`. At most, `buf_size` characters (including the null terminating character) are written to `buf`.

In a similar way as for `printf`, the format string consists of ordinary characters, which are copied unchanged to the `char` array `buf`, and zero or more conversion specifiers. A conversion specifier starts with the character `%` and is followed by a character that indicates the form of transformation required – the supported transformations are given below in Table 3-18. The `strftime` function only supports the "C" locale, and this is reflected in the table.

Table 3-18. Conversion Specifiers Supported by `strftime`

Conversion Specifier	Transformation	ISO/IEC 9899
`%a`	abbreviated weekday name	yes
`%A`	full weekday name	yes
`%b`	abbreviated month name	yes
`%B`	full month name	yes
`%c`	date and time presentation in the form of `DDD MMM dd hh:mm:ss yyyy`	yes

Table 3-18. Conversion Specifiers Supported by `strftime` (Cont'd)

Conversion Specifier	Transformation	ISO/IEC 9899
`%C`	century of the year	POSIX.2-1992 + ISO C99
`%d`	day of the month (`01 - 31`)	yes
`%D`	date represented as `mm/dd/yy`	POSIX.2-1992 + ISO C99
`%e`	day of the month, padded with a space character (cf `%d`)	POSIX.2-1992 + ISO C99
`%F`	date represented as `yyyy-mm-dd`	POSIX.2-1992 + ISO C99
`%h`	abbreviated name of the month (same as `%b`)	POSIX.2-1992 + ISO C99
`%H`	hour of the day as a 24-hour clock (`00-23`)	yes
`%I`	hour of the day as a 12-hour clock (`00-12`)	yes
`%j`	day of the year (`001-366`)	yes
`%k`	hour of the day as a 24-hour clock padded with a space (`0-23`)	no
`%l`	hour of the day as a 12-hour clock padded with a space (`0-12`)	no
`%m`	month of the year (`01-12`)	yes
`%m`	month of the year (`01-12`)	yes
`%M`	minute of the hour (`00-59`)	yes
`%n`	newline character	POSIX.2-1992 + ISO C99
`%p`	`AM` or `PM`	yes
`%P`	`am` or `pm`	no
`%r`	time presented as either `hh:mm:ss AM` or as `hh:mm:ss PM`	POSIX.2-1992 + ISO C99
`%R`	time presented as `hh:mm`	POSIX.2-1992 + ISO C99
`%S`	second of the minute (`00-61`)	yes

Table 3-18. Conversion Specifiers Supported by `strftime` (Cont'd)

Conversion Specifier	Transformation	ISO/IEC 9899
`%t`	tab character	POSIX.2-1992 + ISO C99
`%T`	time formatted as `%H:%M:%S`	POSIX.2-1992 + ISO C99
`%U`	week number of the year (week starts on Sunday) (`00-53`)	yes
`%w`	weekday as a decimal (`0-6`) (`0` if Sunday)	yes
`%W`	week number of the year (week starts on Sunday) (`00-53`)	yes
`%x`	date represented as `mm/dd/yy` (same as `%D`)	yes
`%X`	time represented as `hh:mm:ss`	yes
`%y`	year without the century (`00-99`)	yes
`%Y`	year with the century (`nnnn`)	yes
`%Z`	the time zone name, or nothing if the name cannot be determined	yes
`%%`	`%` character	yes

The current implementation of `time.h` does not support time zones and, therefore, the `%Z` specifier does not generate any characters.

The `strftime` function returns the number of characters (not including the terminating null character) that have been written to `buf`.

Error Conditions

The `strftime` function returns zero if more than `buf_size` characters are required to process the format string. In this case, the contents of the array `buf` will be indeterminate.

Example

```
#include <time.h>
#include <stdio.h>

extern void
print_time(time_t tod)
{
   char tod_string[100];

   strftime(tod_string,
            100,
            "It is %M min and %S secs after %l o'clock (%p)",
            gmtime(&tod));
   puts(tod_string);
}
```

See Also

ctime, gmtime, localtime, mktime

strlen

string length

Synopsis

```
#include <string.h>
size_t strlen(const char *s1);
```

Description

The `strlen` function returns the length of the null-terminated string pointed to by `s1` (not including the terminating null character).

Error Conditions

The `strlen` function does not return an error condition.

Example

```
#include <string.h>
size_t len;

len = strlen("SOMEFUN");      /* len = 7 */
```

See Also

strcspn, strspn

strncat

concatenate characters from one string to another

Synopsis

```
#include <string.h>
char *strncat(char *s1, const char *s2, size_t n);
```

Description

The `strncat` function appends a copy of up to `n` characters in the null-terminated string pointed to by `s2` to the end of the null-terminated string pointed to by `s1`. The function returns a pointer to the new `s1` string.

The behavior of `strncat` is undefined if the two strings overlap. The new `s1` string is terminated with a null character (‘`\0`’).

Error Conditions

The `strncat` function does not return an error condition.

Example

```
#include <string.h>
char string1[50], *ptr;

string1[0]='\0';
strncat(string1, “MOREFUN”, 4);
   /* string1 equals “MORE” */
```

See Also

strcat

strncmp

compare characters in strings

Synopsis

```
#include <string.h>
int strncmp(const char *s1, const char *s2, size_t n);
```

Description

The `strncmp` function lexicographically compares up to `n` characters of the null-terminated strings pointed to by `s1` and `s2`. The function returns a positive value when the `s1` string is greater than the `s2` string, a negative value when the `s2` string is greater than the `s1` string, and a zero when the strings are the same.

Error Conditions

The `strncmp` function does not return an error condition.

Example

```
#include <string.h>
char *ptr1;

ptr1 = "TEST1";
if (strncmp(ptr1, "TEST", 4) == 0)
   printf("%s starts with TEST \n", ptr1);
```

See Also

memcmp, strcmp

strncpy

copy characters from one string to another

Synopsis

```
#include <string.h>
char *strncpy(char *s1, const char *s2, size_t n);
```

Description

The `strncpy` function copies up to `n` characters of the null-terminated string pointed to by `s2` into the space pointed to by `s1`. If the last character copied from `s2` is not a null, the result does not end with a null. The behavior of `strncpy` is undefined when the two objects overlap. The `strncpy` function returns the new `s1`.

If the `s2` string contains fewer than `n` characters, the `s1` string is padded with the null character until all `n` characters are written.

Error Conditions

The `strncpy` function does not return an error condition.

Example

```
#include <string.h>
char string1[50];

strncpy(string1, “MOREFUN”, 4);
                          /* MORE is copied into string1 */
string1[4] = ‘\0’;        /* must null-terminate string1 */
```

See Also

memcpy, memmove, strcpy

strpbrk

find character match in two strings

Synopsis

```
#include <string.h>
char *strpbrk(const char *s1, const char *s2);
```

Description

The `strpbrk` function returns a pointer to the first character in `s1` that is also found in `s2`. The string pointed to by `s2` is treated as a set of characters. The order of the characters in the string is not significant.

Error Conditions

In the event that no character in `s1` matches any in `s2`, a null pointer is returned.

Example

```
#include <string.h>
char *ptr1, *ptr2, *ptr3;

ptr1 = "TESTING";
ptr2 = "SHOP"
ptr3 = strpbrk(ptr1, ptr2);
   /* ptr3 points to the S in TESTING */
```

See Also

strspn

strrchr

find last occurrence of character in string

Synopsis

```
#include <string.h>
char *strrchr(const char *s1, int c);
```

Description

The `strrchr` function returns a pointer to the last occurrence of character `c` in the null-terminated input string `s1`.

Error Conditions

The `strrchr` function returns a null pointer if `c` is not found.

Example

```
#include <string.h>
char *ptr1, *ptr2;

ptr1 = "TESTING";
ptr2 = strrchr(ptr1, 'T');
   /* ptr2 points to the second T of TESTING */
```

See Also

memchr, strchr

strspn

length of segment of characters in both strings

Synopsis

```
#include <string.h>
size_t strspn(const char *s1, const char *s2);
```

Description

The `strspn` function returns the length of the initial segment of `s1` which consists entirely of characters in the string pointed to by `s2`. The string pointed to by `s2` is treated as a set of characters. The order of the characters in the string is not significant.

Error Conditions

The `strspn` function does not return an error condition.

Example

```
#include <string.h>
size_t len;
char *ptr1, *ptr2;

ptr1 = "TESTING";
ptr2 = "ERST";
len = strspn(ptr1, ptr2);      /* len = 4 */
```

See Also

strcspn, strlen

strstr

find string within string

Synopsis

```
#include <string.h>
char *strstr(const char *s1, const char *s2);
```

Description

The `strstr` function returns a pointer to the first occurrence in the string of `s1` of the characters pointed to by `s2`. This excludes the terminating null character in `s1`.

Error Conditions

If the string is not found, `strstr` returns a null pointer. If `s2` points to a string of zero length, `s1` is returned.

Example

```
#include <string.h>
char *ptr1, *ptr2;

ptr1 = "TESTING";
ptr2 = strstr (ptr1, “E”);
   /* ptr2 points to the E in TESTING */
```

See Also

strchr

strtod

convert string to double

Synopsis

```
#include <stdlib.h>
double strtod (const char *nptr, char **endptr)
```

Description

The `strtod` function extracts a value from the string pointed to by `nptr`, and returns the value as a `double`. The `strtod` function expects `nptr` to point to a string that represents either a decimal floating-point number or a hexadecimal floating-point number. Either form of number may be preceded by a sequence of whitespace characters (as determined by the `isspace` function) that the function ignores.

A decimal floating-point number has the form:

```
[sign] [digits] [.digits] [{e|E} [sign] [digits]]
```

The `sign` token is optional and is either plus (+) or minus (–); and digits are one or more decimal digits. The sequence of digits may contain a decimal point (.).

The decimal digits can be followed by an exponent, which consists of an introductory letter (`e` or `E`) and an optionally signed integer. If neither an exponent part nor a decimal point appears, a decimal point is assumed to follow the last digit in the string.

The form of a hexadecimal floating-point number is:

```
[sign] [{0x}|{0X}] [hexdigs] [.hexdigs] [{p|P} [sign] [digits]]
```

A hexadecimal floating-point number may start with an optional plus (+) or minus (–) followed by the hexadecimal prefix `0x` or `0X`. This character sequence must be followed by one or more hexadecimal characters that optionally contain a decimal point (.).

The hexadecimal digits are followed by a binary exponent that consists of the letter `p` or `P`, an optional sign, and a non-empty sequence of decimal digits. The exponent is interpreted as a power of two that is used to scale the fraction represented by the tokens `[hexdigs] [.hexdigs]`.

The first character that does not fit either form of number stops the scan. If `endptr` is not NULL, a pointer to the character that stopped the scan is stored at the location pointed to by `endptr`. If no conversion can be performed, the value of `nptr` is stored at the location pointed to by `endptr`.

Error Conditions

The `strtod` function returns a zero if no conversion is made and a pointer to the invalid string is stored in the object pointed to by `endptr`. If the correct value results in an overflow, a positive or negative (as appropriate) `HUGE_VAL` is returned. If the correct value results in an underflow, zero is returned. The `ERANGE` value is stored in `errno` in the case of either an overflow or underflow.

Example

```
#include <stdlib.h>
char *rem;
double dd;

dd = strtod ("2345.5E4 abc",&rem);
     /* dd = 2.3455E+7, rem = " abc" */

dd = strtod ("-0x1.800p+9,123",&rem);
     /* dd = -768.0, rem = ",123"   */
```

See Also

atof, strtol, strtoul

strtof

convert string to float

Synopsis

```
#include <stdlib.h>
float strtof (const char *nptr, char **endptr)
```

Description

The `strtof` function extracts a value from the string pointed to by `nptr`, and returns the value as a `float`. The `strtof` function expects `nptr` to point to a string that represents either a decimal floating-point number or a hexadecimal floating-point number. Either form of number may be preceded by a sequence of whitespace characters (as determined by the `isspace` function) that the function ignores.

A decimal floating-point number has the form:

```
[sign] [digits] [.digits] [{e|E} [sign] [digits]]
```

The `sign` token is optional and is either plus (+) or minus (–); and digits are one or more decimal digits. The sequence of digits may contain a decimal point (.).

The decimal digits can be followed by an exponent, which consists of an introductory letter (`e` or `E`) and an optionally signed integer. If neither an exponent part nor a decimal point appears, a decimal point is assumed to follow the last digit in the string.

The form of a hexadecimal floating-point number is:

```
[sign] [{0x}|{0X}] [hexdigs] [.hexdigs] [{p|P} [sign] [digits]]
```

A hexadecimal floating-point number may start with an optional plus (+) or minus (–) followed by the hexadecimal prefix `0x` or `0X`. This character sequence must be followed by one or more hexadecimal characters that optionally contain a decimal point (.).

The hexadecimal digits are followed by a binary exponent that consists of the letter `p` or `P`, an optional sign, and a non-empty sequence of decimal digits. The exponent is interpreted as a power of two that is used to scale the fraction represented by the tokens `[hexdigs] [.hexdigs]`.

The first character that does not fit either form of number stops the scan. If `endptr` is not NULL, a pointer to the character that stopped the scan is stored at the location pointed to by `endptr`. If no conversion can be performed, the value of `nptr` is stored at the location pointed to by `endptr`.

Error Conditions

The `strtof` function returns a zero if no conversion is made and a pointer to the invalid string is stored in the object pointed to by `endptr`. If the correct value results in an overflow, a positive or negative (as appropriate) `HUGE_VAL` is returned. If the correct value results in an underflow, zero is returned. The `ERANGE` value is stored in `errno` in the case of either an overflow or underflow.

Example

```
#include <stdlib.h>
char *rem;
float ff;

ff = strtof ("2345.5E4 abc",&rem);
     /* ff = 2.3455E+7, rem = " abc" */

ff = strtof ("-0x1.800p+9,123",&rem);
     /* ff = -768.0, rem = ",123"   */
```

See Also

atof, strtol, strtoul

strtold

convert string to long double

Synopsis

```
#include <stdlib.h>
long double strtold(const char *nptr, char **endptr)
```

Description

The `strtold` function extracts a value from the string pointed to by `nptr`, and returns the value as a `long double`. The `strtold` function expects `nptr` to point to a string that represents either a decimal floating-point number or a hexadecimal floating-point number. Either form of number may be preceded by a sequence of whitespace characters (as determined by the `isspace` function) that the function ignores.

A decimal floating-point number has the form:

```
[sign] [digits] [.digits] [{e|E} [sign] [digits]]
```

The `sign` token is optional and is either plus (+) or minus (–); and digits are one or more decimal digits. The sequence of digits may contain a decimal point (.).

The decimal digits can be followed by an exponent, which consists of an introductory letter (`e` or `E`) and an optionally signed integer. If neither an exponent part nor a decimal point appears, a decimal point is assumed to follow the last digit in the string.

The form of a hexadecimal floating-point number is:

```
[sign] [{0x}|{0X}] [hexdigs] [.hexdigs] [{p|P} [sign] [digits]]
```

A hexadecimal floating-point number may start with an optional plus (+) or minus (–) followed by the hexadecimal prefix `0x` or `0X`. This character sequence must be followed by one or more hexadecimal characters that optionally contain a decimal point (.).

The hexadecimal digits are followed by a binary exponent that consists of the letter `p` or `P`, an optional sign, and a non-empty sequence of decimal digits. The exponent is interpreted as a power of two that is used to scale the fraction represented by the tokens `[hexdigs] [.hexdigs]`.

The first character that does not fit either form of number stops the scan. If `endptr` is not NULL, a pointer to the character that stopped the scan is stored at the location pointed to by `endptr`. If no conversion can be performed, the value of `nptr` is stored at the location pointed to by `endptr`.

Error Conditions

The `strtold` function returns a zero if no conversion can be made and a pointer to the invalid string is stored in the object pointed to by `endptr`. If the correct value results in an overflow, a positive or negative (as appropriate) `LDBL_MAX` is returned. If the correct value results in an underflow, zero is returned. The `ERANGE` value is stored in `errno` in the case of either an overflow or underflow.

Example

```
#include <stdlib.h>
char *rem;
long double dd;

dd = strtold ("2345.5E4 abc",&rem);
     /* dd = 2.3455E+7, rem = " abc" */

dd = strtold ("-0x1.800p+9,123",&rem);
     /* dd = -768.0, rem = ",123"   */
```

See Also

atold, strtol, strtoul

strtok

convert string to tokens

Synopsis

```
#include <string.h>
char *strtok(char *s1, const char *s2);
```

Description

The `strtok` function returns successive tokens from the string `s1`, where each token is delimited by characters from the string `s2`.

A call to `strtok`, with `s1` not NULL, returns a pointer to the first token in `s1`, where a token is a consecutive sequence of characters not in `s2`. The `s1` string is modified in place to insert a null character at the end of the returned token. If `s1` consists entirely of characters from `s2`, NULL is returned.

Subsequent calls to `strtok`, with `s1` equal to NULL, return successive tokens from the same string. When the string contains no further tokens, NULL is returned. Each new call to `strtok` may use a new delimiter string, even if `s1` is NULL. If `s1` is NULL, the remainder of the string is converted into tokens using the new delimiter characters.

Error Conditions

The `strtok` function returns a null pointer if there are no tokens remaining in the string.

Example

```
#include <string.h>
static char str[] = "a phrase to be tested, today";
char *t;
```

```
t = strtok(str, " ");       /* t points to "a"            */
t = strtok(NULL, " ");      /* t points to "phrase"       */
t = strtok(NULL, ",");      /* t points to "to be tested" */
t = strtok(NULL, ".");      /* t points to " today"       */
t = strtok(NULL, ".");      /* t = NULL                   */
```

See Also

No references to this function.

strtol

convert string to long integer

Synopsis

```
#include <stdlib.h>
long int strtol(const char *nptr, char **endptr, int base);
```

Description

The `strtol` function returns as a `long int` the value represented by the string `nptr`. If `endptr` is not a null pointer, `strtol` stores a pointer to the unconverted remainder in `*endptr`.

The `strtol` function breaks down the input into three sections: white space (as determined by `isspace`), initial characters, and unrecognized characters, including a terminating null character. The initial characters may comprise an optional sign character, `0x` or `0X`, when `base` is 16, and those letters and digits which represent an integer with a radix of `base`. The letters (`a-z` or `A-Z`) are assigned the values 10 to 35 and are permitted only when those values are less than the value of `base`.

If `base` is zero, the base is taken from the initial characters. A leading `0x` indicates base 16; a leading `0` indicates base 8. For any other leading characters, base 10 is used. If `base` is between 2 and 36, it is used as a base for conversion.

Error Conditions

The `strtol` function returns a zero if no conversion is made and a pointer to the invalid string is stored in the object pointed to by `endptr` (provided that `endptr` is not a null pointer). If the correct value results in an overflow, positive or negative (as appropriate) `LONG_MAX` is returned. If the correct value results in an underflow, `LONG_MIN` is returned. The `ERANGE` value is stored in `errno` in the case of either overflow or underflow.

Example

```
#include <stdlib.h>
#define base 10
char *rem;
long int i;

i = strtol("2345.5", &rem, base);
    /* i=2345, rem=".5" */
```

See Also

atoi, atol, strtoul

strtoll

convert string to long long integer

Synopsis

```
#include <stdlib.h>
long long int strtoll(const char *nptr, char **endptr, int base);
```

Description

The `strtoll` function returns as a `long long int` the value represented by the string `nptr`. If `endptr` is not a null pointer, `strtoll` stores a pointer to the unconverted remainder in `*endptr`.

The `strtoll` function breaks down the input into three sections: white space (as determined by `isspace`), initial characters, and unrecognized characters, including a terminating null character. The initial characters may comprise an optional sign character, `0x` or `0X`, when `base` is 16, and those letters and digits which represent an integer with a radix of `base`. The letters (`a-z` or `A-Z`) are assigned the values 10 to 35 and are permitted only when those values are less than the value of `base`.

If `base` is zero, the base is taken from the initial characters. A leading `0x` indicates base 16; a leading `0` indicates base 8. For any other leading characters, base 10 is used. If `base` is between 2 and 36, it is used as a base for conversion.

Error Conditions

The `strtoll` function returns a zero if no conversion is made and a pointer to the invalid string is stored in the object pointed to by `endptr` (provided that `endptr` is not a null pointer). If the correct value results in an overflow, positive or negative (as appropriate) `LLONG_MAX` is returned. If the correct value results in an underflow, `LLONG_MIN` is returned. The `ERANGE` value is stored in `errno` in the case of either overflow or underflow.

Example

```
#include <stdlib.h>
#define base 10
char *rem;
long long int i;

i = strtoll("2345.5", &rem, base);
    /* i=2345, rem=".5" */
```

See Also

atoll, strtoull

strtoul

convert string to unsigned long integer

Synopsis

```
#include <stdlib.h>
unsigned long int strtoul(const char *nptr,
         char **endptr, int base);
```

Description

The `strtoul` function returns as an `unsigned long int` the value represented by the string `nptr`. If `endptr` is not a null pointer, `strtoul` stores a pointer to the unconverted remainder in `*endptr`.

The `strtoul` function breaks down the input into three sections:

- whitespace (as determined by `isspace`)
- initial characters
- unrecognized characters including a terminating null character

The initial characters may comprise an optional sign character, `0x` or `0X`, when `base` is 16, and those letters and digits which represent an integer with a radix of `base`. The letters (`a-z` or `A-Z`) are assigned the values 10 to 35 and are permitted only when those values are less than the value of `base`.

If `base` is zero, the base is taken from the initial characters. A leading `0x` indicates base 16; a leading `0` indicates base 8. For any other leading characters, base 10 is used. If `base` is between 2 and 36, it is used as a base for conversion.

Error Conditions

The `strtoul` function returns a zero if no conversion is made and a pointer to the invalid string is stored in the object pointed to by `endptr` (provided that `endptr` is not a null pointer). If the correct value results in an overflow, `ULONG_MAX` is returned. The `ERANGE` value is stored in `errno` in the case of overflow.

Example

```
#include <stdlib.h>
#define base 10

char *rem;
unsigned long int i;

i = strtoul("2345.5", &rem, base);
    /* i = 2345, rem = ".5" */
```

See Also

atoi, atol, strtol

strtoull

convert string to unsigned long long integer

Synopsis

```
#include <stdlib.h>
unsigned long long int strtoull(const char *nptr,
         char **endptr, int base);
```

Description

The `strtoull` function returns as an `unsigned long long int` the value represented by the string `nptr`. If `endptr` is not a null pointer, `strtoull` stores a pointer to the unconverted remainder in `*endptr`.

The `strtoull` function breaks down the input into three sections:

- whitespace (as determined by `isspace`)
- initial characters
- unrecognized characters including a terminating null character

The initial characters may comprise an optional sign character, `0x` or `0X`, when `base` is 16, and those letters and digits which represent an integer with a radix of `base`. The letters (`a-z` or `A-Z`) are assigned the values 10 to 35 and are permitted only when those values are less than the value of `base`.

If `base` is zero, the base is taken from the initial characters. A leading `0x` indicates base 16; a leading `0` indicates base 8. For any other leading characters, base 10 is used. If `base` is between 2 and 36, it is used as a base for conversion.

Error Conditions

The `strtoull` function returns a zero if no conversion is made and a pointer to the invalid string is stored in the object pointed to by `endptr` (provided that `endptr` is not a null pointer). If the correct value results in an overflow, `ULLONG_MAX` is returned. The `ERANGE` value is stored in `errno` in the case of overflow.

Example

```
#include <stdlib.h>
#define base 10

char *rem;
unsigned long long int i;

i = strtoull("2345.5", &rem, base);
    /* i = 2345, rem = ".5" */
```

See Also

atoll, strtoll

strxfrm

transform string using LC_COLLATE

Synopsis

```
#include <string.h>
size_t strxfrm(char *s1, const char *s2, size_t n);
```

Description

The strxfrm function transforms the string pointed to by s2 using the locale-specific category LC_COLLATE. The function places the result in the array pointed to by s1.

If s1 and s2 are transformed and used as arguments to strcmp, the result is identical to the result derived from strcoll using s1 and s2 as arguments. However, since only C locale is implemented, this function does not perform any transformations other than the number of characters. The string stored in the array pointed to by s1 is never more than n characters, including the terminating null character.

The function returns 1. If this value is n or greater, the result stored in the array pointed to by s1 is indeterminate. The s1 can be a null pointer if n is 0.

Error Conditions

The strxfrm function does not return an error condition.

Example

```
#include <string.h>
char string1[50];
strxfrm(string1, "SOMEFUN", 49);
      /* SOMEFUN is copied into string1 */
```

See Also

strcmp, strcoll

tan

tangent

Synopsis

```
#include <math.h>

float tanf (float x);
double tan (double x);
long double tand (long double x);

fract16 tan_fr16 (fract16 x);
```

Description

The tan functions return the tangent of `x`. Both the argument `x` and the function results are in radians. The defined domain for the `tanf` function is [-9099, 9099], and for the `tand` function the domain is [-4.216e8, 4.216e8].

The `tan_fr16` function is defined for fractional input values between [- π/4, π/4]. The output from the function is in the range [-1.0, 1.0).

Error Conditions

The tan functions return a zero if the input argument is not in the defined domain.

Example

```
#include <math.h>

double y;
y = tan (3.14159/4.0); /* y = 1.0 */
```

See Also

atan, atan2

tanh

hyperbolic tangent

Synopsis

```
#include <math.h>

float tanhf (float x);
double tanh (double x);
long double tanhd (long double x);
```

Description

The tanh functions return the hyperbolic tangent of the argument `x`, where `x` is measured in radians.

Error Conditions

The tanh functions do not return an error condition.

Example

```
#include <math.h>
double x, y;
float z, w;

y = tanh (x);
z = tanhf (w);
```

See Also

cosh, sinh

time

calendar time

Synopsis

```
#include <time.h>
time_t time(time_t *t);
```

Description

The `time` function returns the current calendar time which measures the number of seconds that have elapsed since the start of a known epoch. As the calendar time cannot be determined in this implementation of `time.h`, a result of `(time_t) -1` is returned. The function's result is also assigned to its argument, if the pointer to `t` is not a null pointer.

Error Conditions

The `time` function will return the value `(time_t) -1` if the calendar time is not available.

Example

```
#include <time.h>
#include <stdio.h>

if (time(NULL) == (time_t) -1)
    printf("Calendar time is not available\n");
```

See Also

ctime, gmtime, localtime

tolower

convert from uppercase to lowercase

Synopsis

```
#include <ctype.h>
int tolower(int c);
```

Description

The `tolower` function converts the input character to lowercase if it is uppercase; otherwise, it returns the character.

Error Conditions

The `tolower` function does not return an error condition.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   if(isupper(ch))
      printf("tolower=%#04x", tolower(ch));
   putchar('\n');
}
```

See Also

islower, isupper, toupper

toupper

convert from lowercase to uppercase

Synopsis

```
#include <ctype.h>
int toupper(int c);
```

Description

The `toupper` function converts the input character to uppercase if it is in lowercase; otherwise, it returns the character.

Error Conditions

The `toupper` function does not return an error condition.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   if(islower(ch))
      printf("toupper=%#04x", toupper(ch));
   putchar('\n');
}
```

See Also

islower, isupper, tolower

va_arg

get next argument in variable-length list of arguments

Synopsis

```
#include <stdarg.h>
void va_arg(va_list ap, type);
```

Description

The `va_arg` macro is used to walk through the variable-length list of arguments to a function.

After starting to process a variable-length list of arguments with `va_start`, call `va_arg` with the same `va_list` variable to extract arguments from the list. Each call to `va_arg` returns a new argument from the list.

Substitute a `type` name corresponding to the type of the next argument for the `type` parameter in each call to `va_arg`. After processing the list, call `va_end`.

The `stdarg.h` header file defines a pointer type called `va_list` that is used to access the list of variable arguments.

The function calling `va_arg` is responsible for determining the number and types of arguments in the list. The function needs this information to determine how many times to call `va_arg` and what to pass for the type parameter each time. There are several common ways for a function to determine this type of information. The standard C `printf` function reads its first argument looking for % sequences to determine the number and types of its extra arguments. In the example, all of the arguments are of the same type (`char*`), and a termination value (NULL) is used to indicate the end of the argument list. Other methods are also possible.

If a call to `va_arg` is made after all arguments have been processed, or if `va_arg` is called with a type parameter that is different from the type of the next argument in the list, the behavior of `va_arg` is undefined.

Error Conditions

The `va_arg` macro does not return an error condition.

Example

```
#include <stdarg.h>
#include <string.h>
#include <stdlib.h>

char *concat(char *s1,...)
{
   int len = 0;
   char *result;
   char *s;
   va_list ap;

   va_start (ap,s1);
   s = s1;
   while (s){
      len += strlen (s);
      s = va_arg (ap,char *);
   }
   va_end (ap);

   result = malloc (len +7);
   if (!result)
      return result;
   *result = '';
   va_start (ap,s1);
   s = s1;
   while (s){
      strcat (result,s);
      s = va_arg (ap,char *);
   }
   va_end (ap);
   return result;

}
```

See Also

va_start, va_end

va_end

finish processing variable-length list of arguments

Synopsis

```
#include <stdarg.h>
void va_end(va_list ap);
```

Description

The `va_end` macro can only be used after the `va_start` macro has been invoked. A call to `va_end` concludes the processing of a variable length list of arguments that was begun by `va_start`.

Error Conditions

The `va_end` macro does not return an error condition.

Example

See "va_arg" on page 3-227.

See Also

va_arg, va_start

va_start

initialize processing variable-length list of arguments

Synopsis

```
#include <stdarg.h>
void va_start(va_list ap, parmN);
```

Description

The `va_start` macro is used to start processing variable arguments in a function declared to take a variable number of arguments. The first argument to `va_start` should be a variable of type `va_list`, which is used by `va_arg` to walk through the arguments.

The second argument is the name of the last *named* parameter in the function's parameter list; the list of variable arguments immediately follows this parameter. The `va_start` macro must be invoked before either the `va_arg` or `va_end` macro can be invoked.

Error Conditions

The `va_start` macro does not return an error condition.

Example

See "va_arg" on page 3-227.

See Also

va_arg, va_end

4 DSP RUN-TIME LIBRARY

This chapter describes the DSP run-time library which contains a broad collection of functions that are commonly required by signal processing applications. The services provided by the DSP run-time library include support for general-purpose signal processing such as companders, filters, and Fast Fourier Transform (FFT) functions. All these services are Analog Devices extensions to ANSI standard C. These functions are in addition to the C/C++ run-time library functions that are described in Chapter 3, "C/C++ Run-Time Library".

For more information about the algorithms on which many of the DSP run-time library's math functions are based, see Cody, W. J. and W. Waite, *Software Manual for the Elementary Functions,* Englewood Cliffs, New Jersey: Prentice Hall, 1980.

This chapter contains:

- "DSP Run-Time Library Guide" on page 4-2
  contains information about the library and provides a description of the DSP header files that are included with this release of the `ccblkfn` compiler.
- "DSP Run-Time Library Reference" on page 4-43
  contains the complete reference for each DSP run-time library function provided with this release of the `ccblkfn` compiler.

DSP Run-Time Library Guide

The DSP run-time library contains functions that you can call from your source program. This section describes how to use the library and provides information about:

- "Linking DSP Library Functions"
- "Working With Library Source Code" on page 4-3
- "DSP Header Files" on page 4-4
- "Measuring Cycle Counts" on page 4-33

Linking DSP Library Functions

The DSP run-time library is located under the VisualDSP++ installation directory in the subdirectory `Blackfin\lib`. Different versions of the library are supplied and catalogued in Table 4-1.

Table 4-1. DSP Library Files

Blackfin\lib Directory	Description
`libdsp532.dlb` `libdsp535.dlb`	DSP run-time library
`libdsp532y.dlb` `libdsp535y.dlb`	DSP run-time library built with `-workaround all`
`libdspm3res532.dlb` `libdspm3res535.dlb`	DSP run-time library built with `-reserve M3`
`libdspm3res532y.dlb` `libdspm3res535y.dlb`	DSP run-time library built with `-reserve M3 -workaround all`

In general, different versions of the DSP run-time library are supplied in binary form. For instance, one set of libraries have been built for execution on ADSP-BF531, ADSP-BF532, ADSP-BF533, ADSP-BF537, ADSP-BF538, ADSP-BF539 or ADSP-BF561 processors and all the files

have 532 in their filenames, while another set of libraries have been built for execution on ADSP-BF535 or AD6532 processors and all the files have 535 in their filenames.

Versions of the library whose file name end with a y (for example, libdsp532y.dlb) have been built with the compiler's -workaround all switch (on page 1-62) and include all available compiler workarounds for hardware anomalies. Libraries are also supplied that have m3res in their name; these libraries have been built with the compiler's -reserve M3 switch (on page 1-53) and do not use the M3 register (see Emulator documentation).

When an application calls a DSP library function, the call creates a reference that the linker resolves. One way to direct the linker to the library's location is to use the default Linker Description File (<your_target>.ldf). If a customized .LDF file is used to link the application, then the appropriate DSP run-time library should be added to the .LDF file used by the project.

Instead of modifying a customized .LDF file, the -l switch (see "-l library" on page 1-38) can be used to specify which library should be searched by the linker. For example, the -ldsp532 switch adds the library libdsp532.dlb to the list of libraries that the linker examines. For more information on .LDF files, see the *VisualDSP++ 4.0 Linker and Utilities Manual.*

Working With Library Source Code

The source code for the functions in the DSP run-time library is provided with your VisualDSP++ software. By default, the libraries are installed in the directory Blackfin\lib and the source files are copied into the directory Blackfin\lib\src. Each function is kept in a separate file. The file name is the name of the function with .asm or .c extension. If you do not intend to modify any of the run-time library functions, you can delete this directory and its contents to conserve disk space.

The source code is provided so specific functions can be customized as a user requires. To modify these files, proficiency in Blackfin assembly language and an understanding of the run-time environment is needed.

Refer to "C/C++ Run-Time Model and Environment" on page 1-230 for more information.

Before making any modifications to the source code, copy it to a file with a different file name and rename the function itself. Test the function before you use it in your system to verify that it is functionally correct.

Analog Devices only supports the run-time library functions as provided.

DSP Header Files

The DSP header files contains prototypes for all the DSP library functions. When the appropriate `#include` preprocessor command is included in your source, the compiler uses the prototypes to check that each function is called with the correct arguments. The DSP header files included in this release of the `ccblkfn` compiler are:

- "complex.h – Basic Complex Arithmetic Functions"
- "cycle_count.h – Basic Cycle Counting" on page 4-7
- "cycles.h – Cycle Counting with Statistics" on page 4-8
- "filter.h – Filters and Transformations" on page 4-8
- "math.h – Math Functions" on page 4-12
- "matrix.h – Matrix Functions" on page 4-16
- "stats.h – Statistical Functions" on page 4-22
- "vector.h – Vector Functions" on page 4-24
- "window.h – Window Generators" on page 4-31

complex.h – Basic Complex Arithmetic Functions

The `complex.h` header file contains type definitions and basic arithmetic operations for variables of type `complex_float`, `complex_double`, `complex_long_double`, and `complex_fract16`. The complex functions defined in this header file are listed in Table 4-2 on page 4-6.

The following structures are used to represent complex numbers in rectangular coordinates:

```
typedef struct
{
   float re;
   float im;
} complex_float;

typedef struct
{
   double re;
   double im;
} complex_double;

typedef struct
{
   long double re;
   long double im;
} complex_long_double;

typedef struct
{
   fract16 re;
   fract16 im;
} complex_fract16;
```

Details of the basic complex arithmetic functions are included in "DSP Run-Time Library Reference" starting on page 4-43.

Table 4-2. Complex Functions

Description	Prototype

Complex Absolute Value	`double cabs (complex_double a)`
	`float cabsf (complex_float a)`
	`long double cabsd (complex_long_double a)`
	`fract16 cabs_fr16 (complex_fract16 a)`

Complex Addition	`complex_double cadd`
	`(complex_double a, complex_double b)`
	`complex_float caddf`
	`(complex_float a, complex_float b)`
	`complex_long_double caddd`
	`(complex_long_double a, complex_long_double b)`
	`complex_fract16 cadd_fr16`
	`(complex_fract16 a, complex_fract16 b)`

Complex Subtraction	`complex_double csub`
	`(complex_double a, complex_double b)`
	`complex_float csubf`
	`(complex_float a, complex_float b)`
	`complex_long_double csubd`
	`(complex_long_double a, complex_long_double b)`
	`complex_fract16 csub_fr16`
	`(complex_fract16 a, complex_fract16 b)`

Complex Multiply	`complex_double cmlt`
	`(complex_double a, complex_double b)`
	`complex_float cmltf`
	`(complex_float a, complex_float b)`
	`complex_long_double cmltd`
	`(complex_long_double a, complex_long_double b)`
	`complex_fract16 cmlt_fr16`
	`(complex_fract16 a, complex_fract16 b)`

Complex Division	`complex_double cdiv`
	`(complex_double a, complex_double b)`
	`complex_float cdivf`
	`(complex_float a, complex_float b)`
	`complex_long_double cdivd`
	`(complex_long_double a, complex_long_double b)`
	`complex_fract16 cdiv_fr16`
	`(complex_fract16 a, complex_fract16 b)`

Table 4-2. Complex Functions (Cont'd)

Description	Prototype

Get Phase of a Complex Number	`double arg (complex_double a)`
	`float argf (complex_float a)`
	`long double argd (complex_long_double a)`
	`fract16 arg_fr16 (complex_fract16 a)`

Complex Conjugate	`complex_double conj (complex_double a)`
	`complex_float conjf (complex_float a)`
	`complex_long_double conjd (complex_long_double a)`
	`complex_fract16 conj_fr16 (complex_fract16 a)`

Convert Cartesian to Polar Coordinates	`double cartesian (complex_double a, double* phase)`
	`float cartesianf (complex_float a, float* phase)`
	`long double cartesiand`
	`(complex_long_double a, long_double* phase)`
	`fract16 cartesian_fr16 (complex_fract16 a, fract16* phase)`

Convert Polar to Cartesian Coordinates	`complex_double polar`
	`(double mag, double phase)`
	`complex_float polarf`
	`(float mag, float phase)`
	`complex_long_double polard`
	`(long double mag, long double phase)`
	`complex_fract16 polar_fr16`
	`(fract16 mag, fract16 phase)`

Complex Exponential	`complex_double cexp (double a)`
	`complex_long_double cexpd (long double a)`
	`complex_float cexpf (float a)`

Normalization	`complex_double norm (complex_double a)`
	`complex_long_double normd (complex_long_double a)`
	`complex_float normf (complex_float a)`

cycle_count.h – Basic Cycle Counting

The `cycle_count.h` header file provides an inexpensive method for benchmarking C-written source by defining basic facilities for measuring cycle counts. The facilities provided are based upon two macros, and a data type which are described in more detail in the section "Measuring Cycle Counts" on page 4-33.

cycles.h – Cycle Counting with Statistics

The `cycles.h` header file defines a set of five macros and an associated data type that may be used to measure the cycle counts used by a section of C-written source. The macros can record how many times a particular piece of code has been executed and also the minimum, average, and maximum number of cycles used. The facilities that are available via this header file are described in the section "Measuring Cycle Counts" on page 4-33.

filter.h – Filters and Transformations

The `filter.h` header file contains filters used in signal processing. It also includes the A-law and μ-law companders that are used by voice-band compression and expansion applications.

This header file also contains functions that perform key signal processing transformations, including FFTs and convolution.

Various forms of the FFT function are provided by the library corresponding to radix-2, radix-4, and two-dimensional FFTs. The number of points is provided as an argument. The header file also defines a complex FFT function that has been implemented using an optimized radix-4 algorithm. However, this function, `cfftf_fr16`, has certain requirements that may not be appropriate for some applications. The twiddle table for the FFT functions is supplied as a separate argument and is normally calculated once during program initialization.

Library functions are provided to initialize a twiddle table. A table can accommodate several FFTs of different sizes by allocating the table at maximum size, and then using the stride argument of the FFT function to specify the step size through the table. If the stride argument is set to 1, the FFT function uses all the table; if the FFT uses only half the number of points of the largest, the stride is 2.

The functions defined in this header file are listed in Table 4-3 and Table 4-4 and are described in detail in "DSP Run-Time Library Reference" on page 4-43.

Table 4-3. Filter Library

Description	Prototype
Finite Impulse Response Filter	`void fir_fr16 (const fract16 x[], fract16 y[], int n, fir_state_fr16 *s)`
Infinite Impulse Response Filter	`void iir_fr16 (const fract16 x[], fract16 y[], int n, iir_state_fr16 *s)`
Direct Form I Infinite Response Filter	`void iirdf1_fr16 (const fract16 x[ ], fract16 y[ ], int n, iirdf1_fr16_state *s)`
FIR Decimation Filter	`void fir_decima_fr16 (const fract16 x[], fract16 y[], int n, fir_state_fr16 *s)`
FIR Interpolation Filter	`void fir_interp_fr16 (const fract16 x[], fract16 y[], int n, fir_state_fr16 *s)`
Complex Finite Impulse Response Filter	`void cfir_fr16 (const complex_fract16 x[], complex_fract16 y[], int n, cfir_state_fr16 *s)`
Convert Coefficients for DF1 IIR	`void coeff_iirdf1_fr16 (const float acoeff[ ], const float bcoeff[ ], fract16 coeff[ ], int nstages)`

Table 4-4. Transformational Functions

Description	Prototype
Fast Fourier Transforms	
Generate FFT Twiddle Factors	`void twidfft_fr16 (complex_fract16 w[], int n)`

Table 4-4. Transformational Functions (Cont'd)

Description	Prototype
Generate FFT Twiddle Factors for Radix-2 FFT	`void twidfftrad2_fr16 (complex_fract16 w[], int n)`
Generate FFT Twiddle Factors for Radix-4 FFT	`void twidfftrad4_fr16 (complex_fract16 w[], int n)`
Generate FFT Twiddle Factors for 2-D FFT	`void twidfft2d_fr16 (complex_fract16 w[], int n)`
Generate FFT Twiddle Factors for Optimized FFT	`void twidfftf_fr16 (complex_fract16 w[], int n)`
N Point Radix-2 Complex Input FFT	`void cfft_fr16 (const complex_fract16 *in, complex fract16 *t, complex_fract16 *out, const complex_fract16 *w, int wst, int n, int block_exponent, int scale_method)`
N Point Radix-2 Real Input FFT	`void rfft_fr16 (const fract16 *in, complex_fract16 *t, complex_fract16 *out, const complex_fract16 *w, int wst, int n, int block_exponent, int scale_method)`
N Point Radix-2 Inverse FFT	`void ifft_fr16 (const complex_fract16 *in, complex_fract16 *t, complex_fract16 *out, const complex_fract16 *w, int wst, int n, int block_exponent, int scale_method)`
N Point Radix-4 Complex Input FFT	`void cfftrad4_fr16 (const complex_fract16 *in, complex fract16 *t, complex_fract16 *out, const complex_fract16 *w, int wst, int n, int block_exponent, int scale_method)`
N Point Radix-4 Real Input FFT	`void rfftrad4_fr16 (const fract16 *in, complex_fract16 *t, complex_fract16 *out, const complex_fract16 *w, int wst, int n, int block_exponent, int scale_method)`

Table 4-4. Transformational Functions (Cont'd)

Description	Prototype
N Point Radix-4 Inverse Input FFT	`void ifftrad4_fr16 (const complex_fract *in, complex_fract16 *t, complex_fract16 *out, const complex_fract16 *w, int wst, int n, int block_exponent, int scale_method)`
Fast N point Radix-4 Complex Input FFT	`void cfftf_fr16 (const complex_fract16 *in, complex_fract16 *out, const complex_fract16 *w, int wst, int n)`
Nxn Point 2-D Complex Input FFT	`void cfft2d_fr16 (const complex_fract16 *in, complex fract16 *t, complex_fract16 *out, const complex_fract16 *w, int wst, int n, int block_exponent, int scale_method)`
Nxn Point 2-D Real Input FFT	`void rfft2d_fr16 (const fract16 *in, complex_fract16 *t, complex_fract16 *out, const complex_fract16 *w, int wst, int n, int block_exponent, int scale_method)`
Nxn Point 2-D Inverse FFT	`void ifft2d_fr16 (const complex_fract16 *in, complex_fract16 *t, complex_fract16 *out, const complex_fract16 *w, int wst, int n, int block_exponent, int scale_method)`
Convolutions	
Convolution	`void convolve_fr16 (const fract16 cin1[], int clen1, const fract16 cin2[], int clen2, fract16 cout[])`
2-D Convolution	`void conv2d_fr16 (const fract16 *cin1, int crow1, int ccol1, const fract16 *cin2, int crow2, int ccol2, fract16 *cout)`

Table 4-4. Transformational Functions (Cont'd)

Description	Prototype
2-D Convolution 3x3 Matrix	`void conv2d3x3_fr16 (const fract16 *cin, int crow1, int ccol1, const fract16 cin2 [3] [3], fract16 *cout)`
Compression/Expansion	
A-law compression	`void a_compress (const short in[], short out[], int n)`
A-law expansion	`void a_expand (const short in[], short out[], int n)`
µ-law compression	`void mu_compress (const short in[], short out[], int n)`
µ-law expansion	`void mu_expand (const char in[], short out[], int n)`

math.h – Math Functions

The standard math functions have been augmented by implementations for the `float` and `long double` data type, and in some cases, for the `fract16` data type.

Table 4-5 provides a summary of the functions defined by the `math.h` header file. Descriptions of these functions are given under the name of the `double` version in "C Run-Time Library Reference" on page 3-53.

This header file also provides prototypes for a number of additional math functions—`clip`, `copysign`, `max`, and `min`, and an integer function, `countones`. These functions are described in "DSP Run-Time Library Reference" on page 4-43.

Table 4-5. Math Library

Description	Prototype

Absolute Value	`double fabs (double x)`
	`float fabsf (float x)`
	`long double fabsd (long double x)`

Anti-log	`double alog (double x)`
	`float alogf (float x)`
	`long double alogd (long double x)`

Base 10 Anti-log	`double alog10 (double x)`
	`float alog10f (float x)`
	`long double alog10d (long double x)`

Arc Cosine	`double acos (double x)`
	`float acosf (float x)`
	`long double acosd (long double x)`
	`fract16 acos_fr16 (fract16 x)`

Arc Sine	`double asin (double x)`
	`float asinf (float x)`
	`long double asind (long double x)`
	`fract16 asin_fr16 (fract16 x)`

Arc Tangent	`double atan (double x)`
	`float atanf (float x)`
	`long double atand (long double x)`
	`fract16 atan_fr16 (fract16 x)`

Arc Tangent of Quotient	`double atan2 (double x, double y)`
	`float atan2f (float x, float y)`
	`long double atan2d (long double x, long double y)`
	`fract16 atan2_fr16 (fract16 x, fract16 y)`

Ceiling	`double ceil (double x)`
	`float ceilf (float x)`
	`long double ceild (long double x)`

Cosine	`double cos (double x)`
	`float cosf (float x)`
	`long double cosd (long double x)`
	`fract16 cos_fr16 (fract16 x)`

Table 4-5. Math Library (Cont'd)

Description	Prototype

Cotangent	`double cot (double x)`
	`float cotf (float x)`
	`long double cotd (long double x)`

Hyperbolic Cosine	`double cosh (double x)`
	`float coshf (float x)`
	`long double coshd (long double x)`

Exponential	`double exp (double x)`
	`float expf (float x)`
	`long double expd (long double x)`

Floor	`double floor (double x)`
	`float floorf (float x)`
	`long double floord (long double x)`

Floating-Point Remainder	`double fmod (double x, double y)`
	`float fmodf (float x, float y)`
	`long double fmodd (long double x, long double y)`

Get Mantissa and Exponent	`double frexp (double x, int *n)`
	`float frexpf (float x, int *n)`
	`long double frexpd (long double x, int *n)`

Is Not A Number?	`int isnanf (float x)`
	`int isnan (double x)`
	`int isnand (long double x)`

Is Infinity?	`int isinff (float x)`
	`int isinf (double x)`
	`int isinfd (long double x)`

Multiply by Power of 2	`double ldexp(double x, int n)`
	`float ldexpf(float x, int n)`
	`long double ldexpd (long double x, int *n)`

Natural Logarithm	`double log (double x)`
	`float logf (float x)`
	`long double logd (long double x)`

Logarithm Base 10	`double log10 (double x)`
	`float log10f (float x)`
	`long double log10d (long double x)`

Table 4-5. Math Library (Cont'd)

Description	Prototype

Get Fraction and Integer	`double modf (double x, double *i)`
	`float modff (float x, float *i)`
	`long double modfd (long double x, long double *i)`

Power	`double pow (double x, double y)`
	`float powf (float x, float y)`
	`long double powd (long double x, long double y)`

Reciprocal Square Root	`double rsqrt (double x)`
	`float rsqrtf (float x)`
	`long double rsqrtd (long double x)`

Sine	`double sin (double x)`
	`float sinf (float x)`
	`long double sind (long double x)`
	`fract16 sin_fr16 (fract16 x)`

Hyperbolic Sine	`double sinh (double x)`
	`float sinhf (float x)`
	`long double sinhd (long double x)`

Square Root	`double sqrt (double x)`
	`float sqrtf (float x)`
	`long double sqrtd (long double x)`
	`fract16 sqrt_fr16 (fract16 x)`

Tangent	`double tan (double x)`
	`float tanf (float x)`
	`long double tand (long double x)`
	`fract16 tan_fr16 (fract16 x)`

Hyperbolic Tangent	`double tanh (double x)`
	`float tanhf (float x)`
	`long double tanhd (long double x)`

matrix.h – Matrix Functions

The `matrix.h` header file contains matrix functions for operating on real and complex matrices, both matrix-scalar and matrix-matrix operations. See "complex.h – Basic Complex Arithmetic Functions" on page 4-5 for definitions of the complex types.

The matrix functions defined in the `matrix.h` header file are listed in Table 4-6.

Table 4-6. Matrix Functions

Description	Prototype

Real Matrix + Scalar Addition	`void matsadd (const double *a, double b, int n, int m, double *c)`
	`void matsaddf (const float *a, float b, int n, int m, float *c)`
	`void matsaddd (const long double *a, long double b, int n, int m, long double *c)`
	`void matsadd_fr16 (const fract16 *a, fract16 b, int n, int m, fract16 *c)`

Real Matrix – Scalar Subtraction	`void matssub (const double *a, double b, int n, int m, double *c)`
	`void matssubf (const float *a, float b, int n, int m, float *c)`
	`void matssubd (const long double *a, long double b, int n, int m, long double *c)`
	`void matssub_fr16 (const fract16 *a, fract16 b, int n, int m, fract16 *c)`

Table 4-6. Matrix Functions (Cont'd)

Description	Prototype

Real Matrix * Scalar Multiplication	`void matsmlt (const double *a, double b, int n, int m, double *c)`
	`void matsmltf (const float *a, float b, int n, int m, float *c)`
	`void matsmltd (const long double *a, long double b, int n, int m, long double *c)`
	`void matsmlt_fr16 (const fract16 *a, fract16 b, int n, int m, fract16 *c)`

Real Matrix + Matrix Addition	`void matmadd (const double *a, const double *b, int n, int m, double *c)`
	`void matmaddf (const float *a, const float *b, int n, int m, float *c)`
	`void matmaddd (const long double *a, const long double *b, int n, int m, long double *c)`
	`void matmadd_fr16 (const fract16 *a, const fract16 *b, int n, int m, fract16 *c)`

Real Matrix – Matrix Subtraction	`void matmsub (const double *a, const double *b, int n, int m, double *c)`
	`void matmsubf (const float *a, const float *b, int n, int m, float *c)`
	`void matmsubd (const long double *a, const long double *b, int n, int m, long double *c)`
	`void matmsub_fr16 (const fract16 *a, const fract16 *b, int n, int m, fract16 *c)`

Table 4-6. Matrix Functions (Cont'd)

Description	Prototype

Real Matrix * Matrix Multiplication	`void matmmlt (const double *a, int n, int k, const double *b, int m, double *c)`
	`void matmmltf (const float *a, int n, int k, const float *b, int m, float *c)`
	`void matmmltd (const long double *a, int n, int k, const long double *b, int m, long double *c)`
	`void matmmlt_fr16 (const fract16 *a, int n, int k, const fract16 *b, int m, fract16 *c)`

Complex Matrix + Scalar Addition	`void cmatsadd (const complex_double *a, complex_double b, int n, int m, complex_double *c)`
	`void cmatsaddf (const complex_float *a, complex_float b, int n, int m, complex_float *c)`
	`void cmatsaddd (const complex_long_double *a, complex_long_double b, int n, int m, complex_long_double *c)`
	`void cmatsadd_fr16 (const complex_fract16 *a, complex_fract16 b, int n, int m, complex_fract16 *c)`

Table 4-6. Matrix Functions (Cont'd)

Description	Prototype

Complex Matrix – Scalar Subtraction	`void cmatssub`
	`  (const complex_double *a,`
	`  complex_double b,`
	`  int n, int m, complex_double *c)`
	`void cmatssubf`
	`  (const complex_float *a,`
	`  complex_float b,`
	`  int n, int m, complex_float *c)`
	`void cmatssubd`
	`  (const complex_long_double *a,`
	`  complex_long_double b,`
	`  int n, int m, complex_long_double *c)`
	`void cmatssub_fr16`
	`  (const complex_fract16 *a,`
	`  complex_fract16 b,`
	`  int n, int m, complex_fract16 *c)`

Complex Matrix * Scalar Multiplication	`void cmatsmlt`
	`  (const complex_double *a,`
	`  complex_double b,`
	`  int n, int m, complex_double *c)`
	`void cmatsmltf`
	`  (const complex_float *a,`
	`  complex_float b,`
	`  int n, int m, complex_float *c)`
	`void cmatsmltd`
	`  (const complex_long_double *a,`
	`  complex_long_double b,`
	`  int n, int m, complex_long_double *c)`
	`void cmatsmlt_fr16`
	`  (const complex_fract16 *a,`
	`  complex_fract16 b,`
	`  int n, int m, complex_fract16 *c)`

Table 4-6. Matrix Functions (Cont'd)

Description	Prototype

Complex Matrix + Matrix Addition	`void cmatmadd (const complex_double *a, const complex_double *b, int n, int m, complex_double *c)`
	`void cmatmaddf (const complex_float *a, const complex_float *b, int n, int m, complex_float *c)`
	`void cmatmaddd (const complex_long_double *a, const complex_long_double *b, int n, int m, complex_long_double *c)`
	`void cmatmadd_fr16 (const complex_fract16 *a, const complex_fract16 *b, int n, int m, complex_fract16 *c)`

Complex Matrix – Matrix Subtraction	`void cmatmsub (const complex_double *a, const complex_double *b, int n, int m, complex_double *c)`
	`void cmatmsubf (const complex_float *a, const complex_float *b, int n, int m, complex_float *c)`
	`void cmatmsubd (const complex_long_double *a, const complex_long_double *b, int n, int m, complex_long_double *c)`
	`void cmatmsub_fr16 (const complex_fract16 *a, const complex_fract16 *b, int n, int m, complex_fract16 *c)`

Table 4-6. Matrix Functions (Cont'd)

Description	Prototype

Complex Matrix * Matrix Multiplication	`void cmatmmlt (const complex_double *a, int n, int k, const complex_double *b, int m, complex_double *c)`
	`void cmatmmltf (const complex_float *a, int n, int k, const complex_float *b, int m, complex_float *c)`
	`void cmatmmltd (const complex_long_double *a, int n, int k, const complex_long_double *b, int m, complex_long_double *c)`
	`void cmatmmlt_fr16 (const complex_fract16 *a, int n, int k, const complex_fract16 *b, int m, complex_fract16 *c)`

Transpose	`void transpm (const double *a, int n, int m, double *c)`
	`void transpmf (const float *a, int n, int m, float *c)`
	`void transpmd (const long double *a, int n, int m, long double *c)`
	`void transpm_fr16 (const fract16 *a, int n, int m, fract16 *c)`

In most of the function prototypes:

- `*a` is a pointer to input matrix `a [] []`
- `*b` is a pointer to input matrix `b [] []`
- `b` is an input scalar
- `n` is the number of rows
- `m` is the number of columns
- `*c` is a pointer to output matrix `c [] []`

In the `matrix*matrix` functions, `n` and `k` are the dimensions of matrix `a` and `k` and `m` are the dimensions of matrix `b`.

The functions described by this header assume that input array arguments are constant; that is, their contents do not change during the course of the routine. In particular, this means the input arguments do not overlap with any output argument.

stats.h – Statistical Functions

The statistical functions defined in the `stats.h` header file are listed in Table 4-7 and are described in detail in "DSP Run-Time Library Reference" on page 4-43.

Table 4-7. Statistical Functions

Description	Prototype

Autocoherence	`void autocohf`
	`  (const float a[], int n, int m, float c[])`
	`void autocoh`
	`  (const double a[], int n, int m, double c[])`
	`void autocohd`
	`  (const long double a[], int n, int m, long double c[])`
	`void autocoh_fr16`
	`  (const fract16 a[], int n, int m, fract16 c[])`

Autocorrelation	`void autocorrf`
	`  (const float a[], int n, int m, float c[])`
	`void autocorr`
	`  (const double a[], int n, int m, double c[])`
	`void autocorrd`
	`  (const long double a[], int n, int m, long double c[])`
	`void autocorr_fr16`
	`  (const fract16 a[], int n, int m, fract16 c[])`

Table 4-7. Statistical Functions (Cont'd)

Description	Prototype

Cross-coherence	`void crosscohf`
	`  (const float a[], const float b[],`
	`  int n, int m, float c[])`
	`void crosscoh`
	`  (const double a[], const double b[],`
	`  int n, int m, double c[])`
	`void crosscohd`
	`  (const long double a[], const long double b[],`
	`  int n, int m, long double c[])`
	`void crosscoh_fr16`
	`  (const fract16 a[], const fract16 b[],`
	`  int n, int m, fract16 c[])`

Cross-correlation	`void crosscorrf`
	`  (const float a[], const float b[],`
	`  int n, int m, float c[])`
	`void crosscorr`
	`  (const double a[], const double b[],`
	`  int n, int m, double c[])`
	`void crosscorrd`
	`  (const long double a[], const long double b[],`
	`  int n, int m, long double c[])`
	`void crosscorr_fr16`
	`  (const fract16 a[], const fract16 b[],`
	`  int n, int m, fract16 c[])`

Histogram	`void histogramf`
	`  (const float a[], int c[],`
	`  float max, float min, int n, int m)`
	`void histogram`
	`  (const double a[], int c[],`
	`  double max, double min, int n, int m)`
	`void histogramd`
	`  (const long double a[], int c[],`
	`  long double max, long double min, int n, int m)`
	`void histogram_fr16`
	`  (const fract16 a[], int c[],`
	`  fract16 max, fract16 min, int n, int m)`

Table 4-7. Statistical Functions (Cont'd)

Description	Prototype

Mean	`float meanf (const float a[], int n)`
	`double mean (const double a[], int n)`
	`long double meand (const long double a[], int n)`
	`fract16 mean_fr16 (const fract16 a[], int n)`

Root Mean Square	`float rmsf (const float a[], int n)`
	`double rms (const double a[], int n)`
	`long double rmsd (const long double a[], int n)`
	`fract16 rms_fr16 (const fract16 a[], int n)`

Variance	`float varf (const float a[], int n)`
	`double var (const double a[], int n)`
	`long double vard (const long double a[], int n)`
	`fract16 var_fr16 (const fract16 a[], int n)`

Count Zero Crossing	`int zero_crossf (const float a[], int n)`
	`int zero_cross (const double a[], int n)`
	`int zero_crossd (const long double a[], int n)`
	`int zero_cross_fr16 (const fract16 a[], int n)`

vector.h – Vector Functions

The `vector.h` header file contains functions for operating on real and complex vectors, both vector-scalar and vector-vector operations. See "complex.h – Basic Complex Arithmetic Functions" on page 4-5 for definitions of the complex types.The functions in the `vector.h` header file are listed in Table 4-8.

In the **Prototype** column, `a[]` and `b[]` are input vectors, `b` is an input scalar, `c[]` is an output vector, and `n` is the number of elements.The functions assume that input array arguments are constant; that is, their contents will not change during the course of the routine. In particular, this means the input arguments do not overlap with any output argument. In general, better run-time performance is achieved by the vector functions if the input vectors and the output vector are in different memory banks. This structure avoids any potential memory bank collisions.

Table 4-8. Vector Functions

Description	Prototype

Real Vector +
Scalar Addition
	`void vecsadd`
	`  (const double a[], double b,`
	`  double c[], int n)`
	`void vecsaddd`
	`  (const long double a[], long double b,`
	`  long double c[], int n)`
	`void vecsaddf`
	`  (const float a[], float b,`
	`  float c[], int n)`
	`void vecsadd_fr16`
	`  (const fract16 a[], fract16 b,`
	`  fract16 c[], int n)`

Real Vector –
Scalar Subtraction
	`void vecssub`
	`  (const double a[], double b,`
	`  double c[], int n)`
	`void vecssubd`
	`  (const long double a[], long double b,`
	`  long double c[], int n)`
	`void vecssubf`
	`  (const float a[], float b,`
	`  float c[], int n)`
	`void vecssub_fr16`
	`  (const fract16 a[], fract16 b,`
	`  fract16 c[], int n)`

Real Vector *
Scalar Multiplication
	`void vecsmlt`
	`  (const double a[], double b,`
	`  double c[], int n)`
	`void vecsmltd`
	`  (const long double a[], long double b,`
	`  long double c[], int n)`
	`void vecsmltf`
	`  (const float a[], float b,`
	`  float c[], int n)`
	`void vecsmlt_fr16`
	`  (const fract16 a[], fract16 b,`
	`  fract16 c[], int n)`

Table 4-8. Vector Functions (Cont'd)

Description	Prototype

Real Vector +
Vector Addition
	`void vecvadd`
	`  (const double a[], const double b [ ],`
	`  double c[], int n)`
	`void vecvaddd`
	`  (const long double a[], const long double b [ ],`
	`  long double c[], int n)`
	`void vecvaddf`
	`  (const float a[], const float b [ ],`
	`  float c[], int n)`
	`void vecvadd_fr16`
	`  (const fract16 a[], const fract16 b [ ],`
	`  fract16 c[], int n)`

Real Vector –
Vector Subtraction
	`void vecvsub`
	`  (const double a[], const double b [ ],`
	`  double c[], int n)`
	`void vecvsubd`
	`  (const long double a[], const long double b [ ],`
	`  long double c[], int n)`
	`void vecvsubf`
	`  (const float a[], const float b [ ],`
	`  float c[], int n)`
	`void vecvsub_fr16`
	`  (const fract16 a[], const fract16 b [ ],`
	`  fract16 c[], int n)`

Real Vector *
Vector Multiplication
	`void vecvmlt`
	`  (const double a[], const double b  [ ],`
	`  double c[], int n)`
	`void vecvmltd`
	`  (const long double a[], const long double b  [ ],`
	`  long double c[], int n)`
	`void vecvmltf`
	`  (const float a[], const float b [ ],`
	`  float c[], int n)`
	`void vecvmlt_fr16`
	`  (const fract16 a[], const fract16 b [ ],`
	`  fract16 c[], int n)`

Table 4-8. Vector Functions (Cont'd)

Description	Prototype

Maximum Value of Vector Elements	`double vecmax (const double a[], int n)`
	`long double vecmaxd (const long double a[], int n)`
	`float vecmaxf (const float a[], int n)`
	`fract16 vecmax_fr16`
	`  (const fract16 a[], int n)`

Minimum Value of Vector Elements	`double vecmin (const double a[], int n)`
	`long double vecmind (const long double a[], int n)`
	`float vecminf (const float a[], int n)`
	`fract16 vecmin_fr16`
	`  (const fract16 a[], int n)`

Index of Maximum Value of Vector Elements	`int vecmaxloc (const double a[], int n)`
	`int vecmaxlocd (const long double a[], int n)`
	`int vecmaxlocf (const float a[], int n)`
	`int vecmaxloc_fr16`
	`  (const fract16 a[], int n)`

Index of Minimum Value of Vector Elements	`int vecminloc (const double a[], int n)`
	`int vecminlocd (const long double a[], int n)`
	`int vecminlocf (const float a[], int n)`
	`int vecminloc_fr16`
	`  (const fract16 a[], int n)`

Complex Vector + Scalar Addition	`void cvecsadd`
	`  (const complex_double a[],`
	`  complex_double b,`
	`  complex_double c[], int n)`
	`void cvecsaddd`
	`  (const complex_long_double a[],`
	`  complex_long_double b,`
	`  complex_long_double c[], int n)`
	`void cvecsaddf`
	`  (const complex_float a[],`
	`  complex_float b,`
	`  complex_float c[], int n)`
	`void cvecsadd_fr16`
	`  (const complex_fract16 a[],`
	`  complex_fract16 b,`
	`  complex_fract16 c[], int n)`

Table 4-8. Vector Functions (Cont'd)

Description	Prototype

Complex Vector – Scalar Subtraction	`void cvecssub`
	`  (const complex_double a[],`
	`  complex_double b,`
	`  complex_double c[], int n)`
	`void cvecssubd`
	`  (const complex_long_double a[],`
	`  complex_long_double b,`
	`  complex_long_double c[], int n)`
	`void cvecssubf`
	`  (const complex_float a[],`
	`  complex_float b,`
	`  complex_float c[], int n)`
	`void cvecssub_fr16`
	`  (const complex_fract16 a[],`
	`  complex_fract16 b,`
	`  complex_fract16 c[], int n)`

Complex Vector * Scalar Multiplication	`void cvecsmlt(`
	`  (const complex_double a[],`
	`  complex_double b,`
	`  complex_double c[], int n)`
	`void cvecsmltd(`
	`  (const complex_long_double a[],`
	`  complex_long_double b,`
	`  complex_long_double c[], int n)`
	`void cvecsmltf`
	`  (const complex_float a[],`
	`  complex_float b,`
	`  complex_float c[], int n)`
	`void cvecsmlt_fr16`
	`  (const complex_fract16 a[],`
	`  complex_fract16 b,`
	`  complex_fract16 c[], int n)`

Table 4-8. Vector Functions (Cont'd)

Description	Prototype

Complex Vector + Vector Addition	`void cvecvadd`
	`  (const complex_double a[],`
	`  const complex_double b[],`
	`  complex_double c[], int n)`
	`void cvecvaddd`
	`  (const complex_long_double a[],`
	`  const complex_long_double b[],`
	`  complex_long_double c[], int n)`
	`void cvecvaddf`
	`  (const complex_float a[],`
	`  const complex_float b[],`
	`  complex_float c[], int n)`
	`void cvecvadd_fr16`
	`  (const complex_fract16 a[],`
	`  const complex_fract16 b[],`
	`  complex_fract16 c[], int n)`

Complex Vector – Vector Subtraction	`void cvecvsub`
	`  (const complex_double a[],`
	`  const complex_double b[],`
	`  complex_double c[], int n)`
	`void cvecvsubd`
	`  (const complex_long_double a[],`
	`  const complex_long_double b[],`
	`  complex_long_double c[], int n)`
	`void cvecvsubf`
	`  (const complex_float a[],`
	`  const complex_float b[],`
	`  complex_float c[], int n)`
	`void cvecvsub_fr16`
	`  (const complex_fract16 a[],`
	`  const complex_fract16 b[],`
	`  complex_fract16 c[], int n)`

Table 4-8. Vector Functions (Cont'd)

Description	Prototype

Complex Vector * Vector Multiplication	`void cvecvmlt (const complex_double a[], const complex_double b[], complex_double c[], int n)`
	`void cvecvmltd (const complex_long_double a[], const complex_long_double b[], complex_long_double c[], int n)`
	`void cvecvmltf (const complex_float a[], const complex_float b[], complex_float c[], int n)`
	`void cvecvmlt_fr16 (const complex_fract16 a[], const complex_fract16 b[], complex_fract16 c[], int n)`

Table 4-8. Vector Functions (Cont'd)

Description	Prototype

Real Vector Dot Product	`double vecdot`
	`  (const double a[],`
	`  const double b[], int n)`
	`long double vecdotd`
	`  (const long double a[],`
	`  const long double b[], int n)`
	`float vecdotf`
	`  (const float a[],`
	`  const float b[], int n)`
	`fract16 vecdot_fr16`
	`  (const fract16 a[],`
	`  const fract16 b[], int n)`

Complex Vector Dot Product	`complex_double cvecdot`
	`  (const complex_double a[],`
	`  const complex_double b[], int n)`
	`complex_long_double cvecdotd`
	`  (const complex_long_double a[],`
	`  const complex_long_double b[], int n)`
	`complex_float cvecdotf`
	`  (const complex_float a[],`
	`  const complex_float b[], int n)`
	`complex_fract16 cvecdot_fr16`
	`  (const complex_fract16 a[],`
	`  const complex_fract16 b[], int n)`

window.h – Window Generators

The `window.h` header file contains various functions to generate windows based on various methodologies. The functions defined in the `window.h` header file are listed in Table 4-9 and are described in detail in "DSP Run-Time Library Reference" on page 4-43.

For all window functions, a stride parameter `a` can be used to space the window values. The window length parameter `n` equates to the number of elements in the window. Therefore, for a stride `a` of 2 and a length `n` of 10, an array of length 20 is required, where every second entry is untouched.

Table 4-9. Window Generator Functions

Description	Prototype
Generate Bartlett Window	`void gen_bartlett_fr16 (fract16 w[], int a, int n)`
Generate Blackman Window	`void gen_blackman_fr16 (fract16 w[], int a, int n)`
Generate Gaussian Window	`void gen_gaussian_fr16 (fract16 w[], float alpha, int a, int n)`
Generate Hamming Window	`void gen_hamming_fr16 (fract16 w[], int a, int n)`
Generate Hanning Window	`void gen_hanning_fr16 (fract16 w[], int a, int n)`
Generate Harris Window	`void gen_harris_fr16 (fract16 w[], int a, int n)`
Generate Kaiser Window	`void gen_kaiser_fr16 (fract16 w[], float beta, int a, int n)`
Generate Rectangular Window	`void gen_rectangular_fr16 (fract16 w[], int a, int n)`
Generate Triangle Window	`void gen_triangle_fr16 (fract16 w[], int a, int n)`
Generate Vonhann Window	`void gen_vonhann_fr16 (fract16 w[], int a, int n)`

Measuring Cycle Counts

The common basis for benchmarking some arbitrary C-written source is to measure the number of processor cycles that the code uses. Once this figure is known, it can be used to calculate the actual time taken by multiplying the number of processor cycles by the clock rate of the processor. The run-time library provides three alternative methods for measuring processor counts. Each of these methods is described in the following sections:

- "Basic Cycle Counting Facility" on page 4-33
- "Cycle Counting Facility with Statistics" on page 4-35
- "Using time.h to Measure Cycle Counts" on page 4-38
- "Determining the Processor Clock Rate" on page 4-40
- "Considerations when Measuring Cycle Counts" on page 4-41

Basic Cycle Counting Facility

The fundamental approach to measuring the performance of a section of code is to record the current value of the cycle count register before executing the section of code, and then reading the register again after the code has been executed. This process is represented by two macros that are defined in the `cycle_count.h` header file; the macros are:

```
START_CYCLE_COUNT(S)
STOP_CYCLE_COUNT(T,S)
```

The parameter `S` is set by the macro `START_CYCLE_COUNT` to the current value of the cycle count register; this value should then be passed to the macro `STOP_CYCLE_COUNT`, which will calculate the difference between the parameter and current value of the cycle count register. Reading the cycle count register incurs an overhead of a small number of cycles and the

macro ensures that the difference returned (in the parameter `T`) will be adjusted to allow for this additional cost. The parameters `S` and `T` must be separate variables; they should be declared as a `cycle_t` data type which the header file `cycle_count.h` defines as:

```
typedef volatile unsigned long long cycle_t;
```

The header file also defines the macro:

```
PRINT_CYCLES(STRING,T)
```

which is provided mainly as an example of how to print a value of type `cycle_t`; the macro outputs the text `STRING` on `stdout` followed by the number of cycles `T`.

The instrumentation represented by the macros defined in this section is only activated if the program is compiled with the `-DDO_CYCLE_COUNTS` switch. If this switch is not specified, then the macros are replaced by empty statements and have no effect on the program.

The following example demonstrates how the basic cycle counting facility may be used to monitor the performance of a section of code:

```
#include <cycle_count.h>
#include <stdio.h>

extern int
main(void)
{
   cycle_t start_count;
   cycle_t final_count;

   START_CYCLE_COUNT(start_count)
   Some_Function_Or_Code_To_Measure();
   STOP_CYCLE_COUNT(final_count,start_count)

   PRINT_CYCLES("Number of cycles: ",final_count)
}
```

The run-time libraries provide alternative facilities for measuring the performance of C source (see "Cycle Counting Facility with Statistics" on page 4-35 and "Using time.h to Measure Cycle Counts" on page 4-38); the relative benefits of this facility are outlined in "Considerations when Measuring Cycle Counts" on page 4-41.

The basic cycle counting facility is based upon macros; it may therefore be customized for a particular application if required, without the need for rebuilding the run-time libraries.

Cycle Counting Facility with Statistics

The `cycles.h` header file defines a set of macros for measuring the performance of compiled C source. As well as providing the basic facility for reading the cycle count registers of the Blackfin architecture, the macros also have the capability of accumulating statistics that are suited to recording the performance of a section of code that is executed repeatedly.

If the switch `-DDO_CYCLE_COUNTS` is specified at compile-time, then the `cycles.h` header file defines the following macros:

- `CYCLES_INIT(S)`
  a macro that initializes the system timing mechanism and clears the parameter `S`; an application must contain one reference to this macro.

- `CYCLES_START(S)`
  a macro that extracts the current value of the cycle count register and saves it in the parameter `S`.

- `CYCLES_STOP(S)`
  a macro that extracts the current value of the cycle count register and accumulates statistics in the parameter `S`, based on the previous reference to the `CYCLES_START` macro.

- `CYCLES_PRINT(S)`
  a macro which prints a summary of the accumulated statistics recorded in the parameter `S`.

- `CYCLES_RESET(S)`
  a macro which re-zeros the accumulated statistics that are recorded in the parameter `S`.

The parameter `S` that is passed to the macros must be declared to be of the type `cycle_stats_t`; this is a structured data type that is defined in the `cycles.h` header file. The data type has the capability of recording the number of times that an instrumented part of the source has been executed, as well as the minimum, maximum, and average number of cycles that have been used. If an instrumented piece of code has been executed for example, 4 times, the `CYCLES_PRINT` macro would generate output on the standard stream `stdout` in the form:

```
AVG   : 95
MIN   : 92
MAX   : 100
CALLS : 4
```

If an instrumented piece of code had only been executed once, then the `CYCLES_PRINT` macro would print a message of the form:

```
CYCLES : 95
```

If the switch `-DDO_CYCLE_COUNTS` is not specified, then the macros described above are defined as null macros and no cycle count information is gathered. To switch between development and release mode therefore only requires a re-compilation and will not require any changes to the source of an application.

The macros defined in the `cycles.h` header file may be customized for a particular application without the requirement for rebuilding the run-time libraries.

An example that demonstrates how this facility may be used is:

```
#include <cycles.h>
#include <stdio.h>

extern void foo(void);
extern void bar(void);

extern int
main(void)
{
   cycle_stats_t stats;
   int i;

   CYCLES_INIT(stats)

   for (i = 0; i < LIMIT; i++) {
      CYCLES_START(stats)
      foo();
      CYCLES_STOP(stats)
   }
   printf("Cycles used by foo\n");
   CYCLES_PRINT(stats)
   CYCLES_RESET(stats)

   for (i = 0; i < LIMIT; i++) {
      CYCLES_START(stats)
      bar();
      CYCLES_STOP(stats)
   }
   printf("Cycles used by bar\n");
   CYCLES_PRINT(stats)
}
```

This example might output:

```
   Cycles used by foo
      AVG   : 25454
      MIN   : 23003
      MAX   : 26295
      CALLS : 16
```

```
Cycles used by bar
   AVG   : 8727
   MIN   : 7653
   MAX   : 8912
   CALLS : 16
```

Alterative methods of measuring the performance of compiled C source are described in the sections "Basic Cycle Counting Facility" on page 4-33 and "Using time.h to Measure Cycle Counts" on page 4-38. Also refer to "Considerations when Measuring Cycle Counts" on page 4-41 which provides some useful tips with regards to performance measurements.

Using time.h to Measure Cycle Counts

The `time.h` header file defines the data type `clock_t`, the `clock` function, and the macro `CLOCKS_PER_SEC`, which together may be used to calculate the number of seconds spent in a program.

In the ANSI C standard, the `clock` function is defined to return the number of implementation dependent clock "ticks" that have elapsed since the program began, and in this version of the C/C++ compiler the function returns the number of processor cycles that an application has used.

The conventional way of using the facilities of the `time.h` header file to measure the time spent in a program is to call the `clock` function at the start of a program, and then subtract this value from the value returned by a subsequent call to the function. This difference is usually cast to a floating-point type, and is then divided by the macro `CLOCKS_PER_SEC` to determine the time in seconds that has occurred between the two calls.

If this method of timing is used by an application then it is important to note that:

- the value assigned to the macro `CLOCKS_PER_SEC` should be independently verified to ensure that it is correct for the particular processor being used (see "Determining the Processor Clock Rate" on page 4-40),

- the result returned by the `clock` function does not include the overhead of calling the library function.

A typical example that demonstrates the use of the `time.h` header file to measure the amount of time that an application takes is shown below.

```
#include <time.h>
#include <stdio.h>

extern int
main(void)
{
    volatile clock_t clock_start;
    volatile clock_t clock_stop;

    double secs;

    clock_start = clock();
    Some_Function_Or_Code_To_Measure();
    clock_stop = clock();

    secs = ((double) (stop_time - start_time))
           / CLOCKS_PER_SEC;
    printf("Time taken is %e seconds\n",secs);
}
```

The header files `cycles.h` and `cycle_count.h` define other methods for benchmarking an application—these header files are described in the sections "Basic Cycle Counting Facility" on page 4-33 and "Cycle Counting Facility with Statistics" on page 4-35, respectively. Also refer to "Considerations when Measuring Cycle Counts" on page 4-41 which provides some guidelines that may be useful.

Determining the Processor Clock Rate

Applications may be benchmarked with respect to how many processor cycles that they use. However, more typically applications are benchmarked with respect to how much time (for example, in seconds) that they take.

To measure the amount of time that an application takes to run on a Blackfin processor usually involves first determining the number of cycles that the processor takes, and then dividing this value by the processor's clock rate. The `time.h` header file defines the macro `CLOCKS_PER_SEC` as the number of processor "ticks" per second. On Blackfin processors, it is set by the run-time library to one of the following values in descending order of precedence:

- via the compile-time switch `-DCLOCKS_PER_SEC=<definition>`. Because the `time_t` type is based on the `long long int` data type, it is recommended that the value assigned to the symbolic name `CLOCKS_PER_SEC` is defined as the same data type by qualifying the value with the `LL` (or `ll`) suffix (for example, `-DCLOCKS_PER_SEC=60000000LL`).
- via the System Services Library
- via the **Processor speed** box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **Processor** (1) category
- from the **cycles.h** header file

If the value of the macro `CLOCKS_PER_SEC` is taken from the `cycles.h` header file, then be aware that the clock rate of the processor will usually be taken to be the maximum speed of the processor, which is not necessarily the speed of the processor at `RESET`.

Considerations when Measuring Cycle Counts

The run-time library provides three different methods for benchmarking C-compiled code. Each of these alternatives are described in the sections:

- "Basic Cycle Counting Facility" on page 4-33
  The basic cycle counting facility represents an inexpensive and relatively inobtrusive method for benchmarking C-written source using cycle counts. The facility is based on macros that factor-in the overhead incurred by the instrumentation. The macros may be customized and they can be switched either or off, and so no source changes are required when moving between development and release mode. The same set of macros is available on other platforms provided by Analog Devices.

- "Cycle Counting Facility with Statistics" on page 4-35
  This is a cycle-counting facility that has more features than the basic cycle counting facility described above. It is therefore more expensive in terms of program memory, data memory, and cycles consumed. However, it does have the ability to record the number of times that the instrumented code has been executed and to calculate the maximum, minimum, and average cost of each iteration. The macros provided take into account the overhead involved in reading the cycle count register. By default, the macros are switched off, but they are switched on by specifying the `-DDO_CYCLE_COUNTS` compile-time switch. The macros may also be customized for a specific application. This cycle counting facility is also available on other Analog Devices architectures.

- "Using time.h to Measure Cycle Counts" on page 4-38
  The facilities of the `time.h` header file represent a simple method for measuring the performance of an application that is portable across a large number of different architectures and systems. These facilities are based around the `clock` function.

  The `clock` function however does not take into account the cost

involved in invoking the function. In addition, references to the function may affect the code that the optimizer generates in the vicinity of the function call. This method of benchmarking may not accurately reflect the true cost of the code being measured.

This method is more suited to benchmarking applications rather than smaller sections of code that run for a much shorter time span.

When benchmarking code, some thought is required when adding timing instrumentation to C source that will be optimized. If the sequence of statements to be measured is not selected carefully, the optimizer may move instructions into (and out of) the code region and/or it may re-site the instrumentation itself, thus leading to distorted measurements. It is therefore generally considered more reliable to measure the cycle count of calling (and returning from) a function rather than a sequence of statements within a function.

It is recommended that variables that are used directly in benchmarking are simple scalars that are allocated in internal memory (be they assigned the result of a reference to the `clock` function, or be they used as arguments to the cycle counting macros). In the case of variables that are assigned the result of the `clock` function, it is also recommended that they be defined with the `volatile` keyword.

The cycle count registers of the Blackfin architecture are called the `CYCLES` and `CYCLES2` registers. These registers are 32-bit registers. The `CYCLES` register is incremented at every processor cycle; when it wraps back to zero the `CYCLES2` register is incremented. Together these registers represent a 64-bit counter that is unlikely to wrap around to zero during the timing of an application.

DSP Run-Time Library Reference

This section provides descriptions of the DSP run-time library functions.

Notation Conventions
An interval of numbers is indicated by the minimum and maximum, separated by a comma, and enclosed in two square brackets, two parentheses, or one of each. A square bracket indicates that the endpoint is included in the set of numbers; a parenthesis indicates that the endpoint is not included.

Reference Format
Each function in the library has a reference page. These pages have the following format:

Name and Purpose of the function

Synopsis – Required header file and functional prototype; when the functionality is provided for several data types (for example, `float`, `double`, `long double` or `fract16`), several prototypes are given

Description – Function specification

Algorithm – High-level mathematical representation of the function

Domain – Range of values supported by the function

Notes – Miscellaneous information

For some functions, the interface is presented using the "K&R" style for ease of documentation. An ANSI C prototype is provided in the header file.

a_compress

A-law compression

Synopsis

```
#include <filter.h>
void a_compress (in, out, n)
const short in[];   /* Input array                          */
short out[];        /* Output array                         */
int n;              /* Number of elements to be compressed */
```

Description

The `a_compress` function takes a vector of linear 13-bit signed speech samples and performs A-law compression according to ITU recommendation G.711. Each sample is compressed to 8 bits and is returned in the vector pointed to by `out`.

Algorithm

`C(k)=a-law` compression of A(k) for k = 0 to n-1

Domain

Content of input array: –4096 to 4095

a_expand

A-law expansion

Synopsis

```
#include <filter.h>
void a_expand (in, out, n)
const short in[];   /* Input array                        */
short out[];        /* Output array                       */
int n;              /* Number of elements to be expanded */
```

Description

The `a_expand` function inputs a vector of 8-bit compressed speech samples and expands them according to ITU recommendation G.711. Each input value is expanded to a linear 13-bit signed sample in accordance with the A-law definition and is returned in the vector pointed to by `out`.

Algorithm

`C(k)=a-law` expansion of A(*k*) for k = 0 to n-1

Domain

Content of input array: 0 to 255

alog

anti-log

Synopsis

```
#include <math.h>

float alogf (float x);
double alog (double x);
long double alogd (long double x);
```

Description

The alog functions calculate the natural (base `e`) anti-log of their argument. An anti-log function performs the reverse of a `log` function and is therefore equivalent to exponentiation.

The value `HUGE_VAL` is returned if the argument `x` is greater than the function's domain. For input values less than the domain, the functions return 0.0.

Algorithm

$c = e^x$

Domain

`x` = [–87.33 , 88.72] for `alogf()`

`x` = [–708.39 , 709.78] for `alogd()`

Example

```
#include <math.h>

double y;
y = alog(1.0);          /* y = 2.71828... */
```

See Also

alog10, exp, log, pow

alog10

base 10 anti-log

Synopsis

```
#include <math.h>

float alog10f (float x);
double alog10 (double x);
long double alog10d (long double x);
```

Description

The alog10 functions calculate the base 10 anti-log of their argument. An anti-log function performs the reverse of a `log` function and is therefore equivalent to exponentiation. Therefore, `alog10(x)` is equivalent to `exp(x * log(10.0))`.

The value `HUGE_VAL` is returned if the argument `x` is greater than the function's domain. For input values less than the domain, the functions return 0.0.

Algorithm

$c = e^{(x * \log(10.0))}$

Domain

`x` = [–37.92 , 38.53] for `alog10f()`

`x` = [–307.65 , 308.25] for `alog10d()`

Example

```
#include <math.h>
```

```
double y;
y = alog10(1.0);          /* y = 10.0 */
```

See Also

alog, exp, log10, pow

arg

get phase of a complex number

Synopsis

```
#include <complex.h>
float argf (complex_float a);
double arg (complex_double a);
long double argd (complex_long_double a);
fract16 arg_fr16 (complex_fract16 a);
```

Description

These functions compute the phase associated with a Cartesian number, represented by the complex argument `a`, and return the result.

Refer to the description of the `polar_fr16` function which explains how a phase, represented as a fractional number, is interpreted in polar notation (see "polar" on page 4-131).

Algorithm

$$c = atan\left(\frac{\mathrm{Im}(a)}{\mathrm{Re}(a)}\right)$$

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `argf( )`
-1.7×10^{308} to $+1.7 \times 10^{308}$	for `argd( )`
–1.0 to +1.0	for `arg_fr16( )`

Note

Im (a) /Re (a) < =1	for `arg_fr16 ( )`

autocoh

autocoherence

Synopsis

```
#include <stats.h>

void autocohf (const float in[ ],
               int samples, int lags, float out[ ]);

void autocoh (const double in[ ],
               int samples, int lags, double out[ ]);

void autocohd (const long double in[ ],
               int samples, int lags, long double out[ ]);

void autocoh_fr16 (const fract16 in[ ],
               int samples, int lags, fract16 out[ ]);
```

Description

The autocoh functions compute the autocoherence of the input vector `in[]`, which contain `samples` values. The autocoherence of an input signal is its autocorrelation minus its mean squared. The functions return the result in the output array `out[]` of length `lags`.

Algorithm

$$c_k = \frac{1}{n} * \left(\sum_{j=0}^{n-k-1} (a_j - \bar{a}) * (a_{j+k} - \bar{a}) \right)$$

where k = {0,1,...,m-1}, `m` is the number of `lags`, `n` is the size of the input vector `a`, and $\bar{a}$ is the mean value of input vector `a`.

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `autocohf( )`
-1.7×10^{308} to $+1.7 \times 10^{308}$	for `autocohd( )`
-1.0 to 1.0	for `autocoh_fr16( )`

autocorr

autocorrelation

Synopsis

```
#include <stats.h>

void autocorrf (const float in[ ],
                int samples, int lags, float out[ ]);

void autocorr (const double in[ ],
                int samples, int lags, double out[ ]);

void autocorrd (const long double in[ ],
                int samples, int lags, long double out[ ]);

void autocorr_fr16 (const fract16 in[ ],
                int samples, int lags, fract16 out[ ]);
```

Description

The autocorr functions perform an autocorrelation of a signal. Autocorrelation is the cross-correlation of a signal with a copy of itself. It provides information about the time variation of the signal. The signal to be autocorrelated is given by the `in[]` input array. The number of samples of the autocorrelation sequence to be produced is given by `lags`. The length of the input sequence is given by `samples`.

Autocorrelation is used in digital signal processing applications such as speech analysis..

Algorithm

$$c_k = \frac{1}{n} * \left(\sum_{j=0}^{n-k-1} a_j * a_{j+k} \right)$$

where k = {0,1,...,m-1}, `m` is the number of `lags`, and `n` is the size of the input vector `a`.

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `autocorrf( )`
-1.7×10^{308} to $+1.7 \times 10^{308}$	for `autocorrd( )`
–1.0 to + 1.0	for `autocorr_fr16( )`

cabs

complex absolute value

Synopsis

```
#include <complex.h>
float cabsf (complex_float a);
double cabs (complex_double a);
long double cabsd (complex_long_double a);
fract16 cabs_fr16 (fract16 a);
```

Description

The cabs functions compute the complex absolute value of a complex input and return the result.

Algorithm

$$c = \sqrt{\mathrm{Re}^2(a) + \mathrm{Im}^2(a)}$$

Domain

$\mathrm{Re}^2(a) + \mathrm{Im}^2(a) <= 3.4 \times 10^{38}$ for `cabsf( )`

$\mathrm{Re}^2(a) + \mathrm{Im}^2(a) <= -1.7 \times 10^{308}$ for `cabsd( )`

$\mathrm{Re}^2(a) + \mathrm{Im}^2(a) <= 1.0$ for `cabs_fr16( )`

Note

The minimum input value for both real and imaginary parts can be less than 1/256 for `cabs_fr16` but the result may have bit error of 2–3 bits.

cadd

complex addition

Synopsis

```
#include <complex.h>
complex_float caddf (complex_float a, complex_float b);
complex_double cadd (complex_double a, complex_double b);
complex_long_double caddd (complex_long_double a,
                           complex_long_double b);
complex_fract16 cadd_fr16 (complex_fract16 a, complex_fract16 b);
```

Description

The cadd functions compute the complex addition of two complex inputs, `a` and `b`, and return the result.

Algorithm

```
Re(c) = Re(a) + Re(b)
Im(c) = Im(a) + Im(b)
```

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$ for `caddf( )`

-1.7×10^{308} to $+1.7 \times 10^{308}$ for `caddd( )`

–1.0 to +1.0 for `cadd_fr16( )`

cartesian

convert Cartesian to polar notation

Synopsis

```
#include <complex.h>

float cartesianf (complex_float a, float *phase);
double cartesian (complex_double a, double *phase);
long double cartesiand (complex_long_double a,
                        long double *phase);
fract16 cartesian_fr16 (complex_fract16 a, fract16 *phase);
```

Description

The cartesian functions transform a complex number from Cartesian notation to polar notation. The Cartesian number is represented by the argument `a` that the function converts into a corresponding magnitude, which it returns as the function's result, and a phase that is returned via the second argument `phase`.

Refer to the description of the `polar_fr16` function which explains how a phase, represented as a fractional number, is interpreted in polar notation (see "polar" on page 4-131).

Algorithm

```
magnitude = cabs(a)
phase = arg(a)
```

Domain

-3.4 x 10^{38} to +3.4 x 10^{38}	for `cartesianf( )`
–1.7 x 10^{308} to +1.7 x 10^{308}	for `cartesiand( )`
-1.0 to +1.0	for `cartesian_fr16( )`

Example

```
#include <complex.h>

complex_float point = {-2.0 , 0.0};
float phase;
float mag;
mag = cartesianf (point,&phase);   /* mag = 2.0, phase = π */
```

cdiv

complex division

Synopsis

```
#include <complex.h>
complex_float cdivf (complex_float a, complex_float b);
complex_double cdiv (complex_double a, complex_double b);
complex_long_double cdivd (complex_long_double a,
                           complex_long_double b);
complex_fract16 cdiv_fr16 (complex_fract16 a, complex_fract16 b);
```

Description

The cdiv functions compute the complex division of complex input `a` by complex input `b`, and return the result.

Algorithm

$$\mathrm{Re}(c) = \frac{\mathrm{Re}(a) * \mathrm{Re}(b) + \mathrm{Im}(a) * \mathrm{Im}(b)}{\mathrm{Re}^2(b) + \mathrm{Im}^2(b)}$$

$$\mathrm{Im}(c) = \frac{\mathrm{Re}(b) * \mathrm{Im}(a) - \mathrm{Im}(b) * \mathrm{Re}(a)}{\mathrm{Re}^2(b) + \mathrm{Im}^2(b)}$$

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$ for `cdivf( )`

-1.7×10^{308} to $+1.7 \times 10^{308}$ for `cdivd( )`

–1.0 to 1.0 for `cdiv_fr16( )`

cexp

complex exponential

Synopsis

```
#include <complex.h>
complex_float cexpf (float a);
complex_double cexp (double a);
complex_long_double cexp (long double a);
```

Description

The cexp functions compute the complex exponential of real input `a` and return the result.

Algorithm

```
Re(c) = cos(a)
Im(c) = sin(a)
```

Domain

a = [–9099 ... 9099] for `cexpf ( )`

a = [-8.433e8 ... 8.433e8] for `cexpd ( )`

cfft

n point radix-2 complex input FFT

Synopsis

```
#include <filter.h>
void cfft_fr16(in[], t[], out[], w[], wst, n, block_exponent,
               scale_method)
const complex_fract16 in[];   /* input sequence               */
complex_fract16 t[];          /* temporary working buffer     */
complex_fract16 out[];        /* output sequence              */
const complex_fract16 w[]     /* twiddle sequence             */
int wst;                      /* twiddle factor stride        */
int n;                        /* number of FFT points         */
int block_exponent;           /* block exponent of output data */
int scale_method;             /* scaling method desired:
                                 0-none, 1-static, 2-dynamic  */
```

Description

This function transforms the time domain complex input signal sequence to the frequency domain by using the radix-2 Fast Fourier Transform (FFT).

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n`, where `n` represents the number of points in the FFT. By allocating these arrays in different memory banks, any potential data bank collisions are avoided, thus improving run-time performance. If the input data can be overwritten, the optimum memory usage can be achieved by also specifying the input array as the output array.

The twiddle table is passed in the argument `w`, which must contain at least `n/2` twiddle factors. The function `twidfftrad2_fr16` may be used to initialize the array. If the twiddle table contains more factors than needed for a particular call on `cfft_fr16`, then the stride factor has to be set appropriately; otherwise it should be set to 1.

The arguments `block_exponent` and `scale_method` have been added for future expansion. These arguments are ignored by the function. To avoid overflow, the function scales the output by `1/n`.

Algorithm

$$X(k) = \sum_{n=0}^{N-1} x(n) W_N^{nk}$$

When the sequence length, `n`, equals power of 4, the `cfftrad4` algorithm is also available.

Domain

Input sequence length `n` must be a power of 2 and at least 16.

cfftf

fast N-point radix-4 complex input FFT

Synopsis

```
#include <filter.h>
void cfftf_fr16(in[],out[],twid[],wst,n)
const complex_fract16 in[];     /* input sequence            */
complex_fract16 out[];          /* output sequence           */
const complex_fract16 twid[]    /* twiddle sequence          */
int wst;                        /* twiddle factor stride     */
int n;                          /* number of FFT points      */
```

Description

The `cfftf_fr16` function transforms the time domain complex input signal sequence to the frequency domain by using the accelerated version of the "Discrete Fourier Transform" known as a "Fast Fourier Transform" or FFT. It "decimates in frequency" using an optimized radix-4 algorithm.

The size of the input array `in` and the output array `out` is `n` where `n` represents the number of points in the FFT. The `cfftf_fr16` function has been designed for optimum performance and requires that the input array `in` be aligned on an address boundary that is a multiple of four times the FFT size. For certain applications, this alignment constraint may not be appropriate; in such cases, the application should call the `cfftrad4_fr16` function ("cfftrad4" on page 4-66) instead, with no loss of facility (apart from performance).

The number of points in the FFT, `n`, must be a power of 4 and must be at least 16.

The twiddle table is passed in the argument `twid`, which must contain at least `3/4n` complex twiddle factors. The table should be initialized with complex twiddle factors in which the real coefficients are positive cosine values and the imaginary coefficients are negative sine values. The func-

tion `twidfftf_fr16` (see on page 4-146) may be used to initialize the array. If the twiddle table contains more factors than required for a particular FFT size, then the stride factor `wst` has to be set appropriately; otherwise it should be set to 1.

It is recommended that the output array not be allocated in the same 4K memory sub-bank as either the input array or the twiddle table, as the performance of the function may otherwise degrade due to data bank collisions.

The function uses static scaling of intermediate results to prevent overflow and the final output therefore is scaled by `1/n`.

Versions of this function that do not use the `M3` register are not available. (The `M3` register may be used by an emulator for context switching. Refer to the appropriate Emulator documentation)

Algorithm

$$X(k) = \sum_{n=0}^{N-1} x(n) W_N^{nk}$$

The `cfft_fr16` function (see "cfft" on page 4-61), which uses a radix-2 algorithm, must be used when the FFT size, `n`, is only a power of 2.

Domain

The number of points in the FFT must be a power of 4 and must be at least 16.

Example

```
#include <filter.h>

#define FFTSIZE 64
```

```
#pragma align 256
segment ("seg_1") complex_fract16 input[FFTSIZE];

#pragma align 4
segment ("seg_2") complex_fract16 output[FFTSIZE];

#pragma align 4
segment ("seg_3") complex_fract16 twid[(3*FFTSIZE)/4];

twidfftf_fr16(twid,FFTSIZE);
cfftf_fr16(input,
           output,
           twid,1,FFTSIZE);
```

cfftrad4

n point radix-4 complex input FFT

Synopsis

```
#include <filter.h>
void cfftrad4_fr16 (in[], t[], out[], w[], wst, n,
                    block_exponent, scale_method)
const complex_fract16 in[];   /* input sequence               */
complex_fract16 t[];          /* temporary working buffer     */
complex_fract16 out[];        /* output sequence              */
const complex_fract16 w[]     /* twiddle sequence             */
int wst;                      /* twiddle factor stride        */
int n;                        /* number of FFT points         */
int block_exponent;           /* block exponent of output data */
int scale_method;             /* scaling method desired:
                                 0-none, 1-static, 2-dynamic  */
```

Description

This function transforms the time domain complex input signal sequence to the frequency domain by using the radix-4 Fast Fourier Transform. The `cfftrad4_fr16` function "decimates in frequency" by the radix-4 FFT algorithm.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n`, where `n` represents the number of points in the FFT. Memory bank collisions, which have an adverse effect on run-time performance, may be avoided by allocating all input and working buffers to different memory banks. If the input data can be overwritten, the optimum memory usage can be achieved by also specifying the input array as the output array.

The twiddle table is passed in the argument `w`, which must contain at least `3/4n` twiddle coefficients. The function `twidfftrad4_fr16` may be used to initialize the array. If the twiddle table contains more coefficients than needed for a particular call on `cfftrad4_fr16`, then the stride factor has to be set appropriately; otherwise it should be set to 1.

The arguments `block_exponent` and `scale_method` have been added for future expansion. These arguments are ignored by the function. To avoid overflow, the function performs static scaling by dividing the input by `n`.

Algorithm

$$X(k) = \sum_{n=0}^{N-1} x(n) W_N^{nk}$$

When the sequence length, `n`, is not a power of 4, the radix-2 method must be used. See "cfft" on page 4-61.

Domain

Input sequence length `n` must be a power of 4 and at least 16.

cfft2d

n x n point 2-D complex input FFT

Synopsis

```
#include <filter.h>
void cfft2d_fr16(*in, *t, *out, w[], wst, n, block_exponent,
                 scale_method)
const complex_fract16 *in;  /* pointer to input matrix a[n][n] */
complex_fract16 *t          /* pointer to working buffer t[n][n] */
complex_fract16 *out;       /* pointer to output matrix c[n][n] */
const complex_fract16 w[];  /* twiddle sequence                 */
int wst;                    /* twiddle factor stride            */
int n;                      /* number of FFT points             */
int block_exponent;         /* block exponent of output data    */
int scale_method;           /* scaling method desired:
                               0-none, 1-static, 2-dynamic    */
```

Description

This function computes the two-dimensional Fast Fourier Transform (FFT) of the complex input matrix `a[n][n]` and stores the result to the complex output matrix `c[n][n]`.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n*n`, where `n` represents the number of points in the FFT. Memory bank collisions, which have an adverse effect on run-time performance, may be avoided by allocating all input and working buffers to different memory banks. If the input data can be overwritten, the optimum memory usage can be achieved by also specifying the input array as the output array.

The twiddle table is passed in the argument `w`, which must contain at least `n` twiddle factors. The function `twidfft2d_fr16` may be used to initialize the array. If the twiddle table contains more factors than needed for a particular call on `cfft2d_fr16`, then the stride factor has to be set appropriately; otherwise it should be set to 1.

The arguments `block_exponent` and `scale_method` have been added for future expansion. These arguments are ignored by the function. To avoid overflow, the function scales the output by `n*n`.

Algorithm

$$c(i,j) = \sum_{k=0}^{n-1}\sum_{l=0}^{n-1} a(k,l) * e^{-2\pi j(i*k+j*l)/n}$$

where i={0,1,...,n-1}, j={0,1,2,...,n-1}

Domain

Input sequence length `n` must be a power of 2 and at least 16.

cfir

complex finite impulse response filter

Synopsis

```
#include <filter.h>

void cfir_fr16(x,y,n,s)
const complex_fract16 x[];    /* Input sample vector x    */
complex_fract16 y[];          /* Output sample vector y   */
int n;                        /* Number of input samples  */
cfir_state_fr16 *s;           /* Pointer to filter state
                                 structure                */
```

The function uses the following structure to maintain the state of the filter.

```
typedef struct
{
   int k;                    /* Number of coefficients  */
   complex_fract16 *h;       /* Filter coefficients     */
   complex_fract16 *d;       /* Start of delay line     */
   complex_fract16 *p;       /* Read/write pointer      */
}  cfir_state_fr16;
```

Description

The `cfir_fr16` function implements a complex finite impulse response (CFIR) filter. It generates the filtered response of the complex input data `x` and stores the result in the complex output vector `y`.

The function maintains the filter state in the structured variable `s`, which must be declared and initialized before calling the function. The macro `cfir_init`, in the `filter.h` header file, is available to initialize the structure.

It is defined as:

```
#define cfir_init(state, coeffs, delay, ncoeffs) \
   (state).h = (coeffs);  \
   (state).d = (delay);   \
   (state).p = (delay);   \
   (state).k = (ncoeffs)
```

The characteristics of the filter (passband, stopband, and so on) are dependent upon the number of complex filter coefficients and their values. A pointer to the coefficients should be stored in `s->h`, and `s->k` should be set to the number of coefficients.

Each filter should have its own delay line which is a vector of type `complex_fract16` and whose length is equal to the number of coefficients. The vector should be cleared to zero before calling the function for the first time and should not otherwise be modified by the user program. The structure member `s->d` should be set to the start of the delay line, and the function uses `s->p` to keep track of its current position within the vector.

Algorithm

$$y(k) = \sum_{i=0}^{p-1} h(i) * x(k-i) \text{ for } k = 0,1,\ldots n$$

Domain

–1.0 to +1.0

clip

clip

Synopsis

```
#include <math.h>

int clip (int parm1, int parm2);
long int lclip (long int parm1, long int parm2);
long long int llclip (long long int parm1,
                      long long int parm2);

float fclipf (float parm1, float parm2);
double fclip (double parm1, double parm2);
long double fclipd (long double parm1, long double parm2);

fract16 clip_fr16 (fract16 parm1, fract16 parm2);
```

Description

The clip functions return the first argument if it is less than the absolute value of the second argument; otherwise they return the absolute value of the second argument if the first is positive, or minus the absolute value if the first argument is negative.

Algorithm

```
If ( |parm1| < |parm2|)
    return( parm1)
else
    return( |parm2| * signof(parm1))
```

Domain

Full range for various input parameter types.

cmlt

complex multiply

Synopsis

```
#include <complex.h>
complex_float cmltf (complex_float a, complex_float b)
complex_double cmlt (complex_double a, complex_double b)
complex_long_double cmltd (complex_long_double a,
                           complex_long_double b)
complex_fract16 cmlt_fr16 (complex_fract16 a, complex_fract16 b)
```

Description

The cmlt functions compute the complex multiplication of two complex inputs, `a` and `b`, and return the result.

Algorithm

```
Re(c) = Re(a) * Re(b) - Im(a) * Im(b)
Im(c) = Re(a) * Im(b) + Im(a) * Re(b)
```

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `cmltf( )`
-1.7×10^{308} to $+1.7 \times 10^{308}$	for `cmltd( )`
–1.0 to 1.0	for `cmlt_fr16( )`

coeff_iirdf1

convert coefficients for DF1 IIR filter

Synopsis

```
#include <filter.h>
void coeff_iirdf1_fr16 (const float acoeff[ ],
                        const float bcoeff[ ],
                        fract16 coeff[ ], int nstages);
```

Description

The `coeff_iirdf1_fr16` function transforms a set of A-coefficients and a set of B-coefficients into a set of coefficients for the `iirdf1_fr16` function (see on page 4-121), which implements an optimized, direct form I infinite impulse response (IIR) filter.

The `A-`coefficients and the `B-`coefficients are passed into the function via the floating-point vectors `acoeff` and `bcoeff`, respectively. The `A0` coefficients are assumed to be 1.0, and all other `A-`coefficients must be scaled according; the `A0` coefficients should not be included in the vector `acoeffs`. The number of stages in the filter is given by the parameter `nstages`, and therefore the size of the `acoeffs` vector is `2*nstages` and the size of the `bcoeffs` vector is `(2*nstages) + 1`.

The `coeff_iirdf1_fr16` function scales the coefficients and stores them in the vector `coeff`. The function also stores the appropriate scaling factor in the vector which the `iirdf1_fr16` function will then apply to the filtered response that it generates (thus eliminating the need to scale the output generated by the IIR function). The size of `coeffs` array should be `(4*nstages) + 2`.

Algorithm

The `A`-coefficients and the `B`-coefficients represent the numerator and denominator coefficients of `H(z)`, where `H(z)` is defined as:

$$H(z) = \frac{B(z)}{A(z)} = \frac{b_1 + b_2 z^{-1} + \ldots + b_{m+1} z^{-m}}{a_1 + a_2 z^{-1} + \ldots + a_{m+1} z^{-m}}$$

If any of the coefficients are greater than `0.999969` (the largest floating-point value that can be converted to a value of type `fract16`), then all the `A`-coefficients and all the `B`-coefficients are scaled to be less than 1.0. The coefficients are stored into the vector `coeffs` in the following order:

$[\, b_0 , -a_0 1 , b_0 1 , -a_0 2, b_0 2, \ldots ,$
$-a_n 1 , b_n 1 , -a_n 2 , b_n 2 , \text{scale factor}]$

where `n` is the number of stages.

Note that the `A`-coefficients are negated by the function.

conj

complex conjugate

Synopsis

```
#include <complex.h>
complex_float conjf (complex_float a)
complex_double conj (complex_double a)
complex_long_double conjd (complex_long_double a)
complex_fract16 conj_fr16 (complex_fract16 a)
```

Description

The conj functions conjugate the complex input `a` and return the result.

Algorithm

```
Re(c) = Re(a)
Im(c) = -Im(a)
```

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$ for `conjf( )`,

-1.7×10^{308} to $+1.7 \times 10^{308}$ for `conjd( )`

–1.0 to 1.0 for `conj_fr16( )`

convolve

convolution

Synopsis

```
#include <filter.h>
void convolve_fr16(cin1, clen1, cin2, clen2, cout)
const fract16 cin1[];    /* pointer to input sequence 1    */
int clen1;               /* length of the input sequence 1 */
const fract16 cin2[];    /* pointer to input sequence 2    */
int clen2;               /* length of the input sequence 2 */
fract16 cout[];          /* pointer to output sequence     */
```

Description

This function convolves two sequences pointed to by `cin1` and `cin2`. If `cin1` points to the sequence whose length is `clen1` and `cin2` points to the sequence whose length is `clen2`, the resulting sequence pointed to by `cout` has length `clen1 + clen2 – 1`.

Algorithm

Convolution between two sequences `cin1` and `cin2` is described as:

$$cout[n] = \sum_{k=0}^{clen2-1} cin1[n+k-(clen2-1)] \bullet cin2[(clen2-1)-k]$$

for `n = 0` to `clen1 + clen2-2`.
(Values for `cin1[j]` are considered to be zero for `j < 0` or `j > clen1-1`).

Example

Here is an example of a convolution where `cin1` is of length 4 and `cin2` is of length 3. If we represent `cin1` as "A" and `cin2` as "B", the elements of the output vector are:

```
{A[0]*B[0],

  A[1]*B[0] +  A[0]*B[1],

  A[2]*B[0] +  A[1]*B[1] + A[0]*B[2],

  A[3]*B[0] +  A[2]*B[1] + A[1]*B[2],

               A[3]*B[1] + A[2]*B[2],

                           A[3]*B[2]}
```

Domain

–1.0 to +1.0

conv2d

2-D convolution

Synopsis

```
#include <filter.h>
void conv2d_fr16(min1, mrow1, mcol1, min2, mrow2, mcol2, mout )
const fract16 *min1;      /* pointer to input matrix 1     */
int mrow1;                /* number of rows in matrix 1    */
int mcol1;                /* number of columns in matrix 1 */
const fract16 *min2;      /* pointer to input matrix 2     */
int mrow2;                /* number of rows in matrix 2    */
int mcol2;                /* number of columns in matrix 2 */
fract16 *mout;            /* pointer to output matrix      */
```

Description

The `conv2d` function computes the two-dimensional convolution of input matrix `min1` of size `mrow1*mcol1` and `min2` of size `mrow2*mcol2` and stores the result in matrix `mout` of dimension (`mrow1 + mrow2-1`) `x` (`mcol1 + mcol2-1`).

A temporary work area is allocated from the run-time stack that the function uses to preserve accuracy while evaluating the algorithm. The stack may therefore overflow if the sizes of the input matrices are sufficiently large. The size of the stack may be adjusted by making appropriate changes to the `.LDF` file

Algorithm

Two dimensional input matrix `min1` is convolved with input matrix `min2`, placing the result in a matrix pointed to by the `mout`.

$$mout[c, r] = \sum_{i=0}^{mcol2-1} \sum_{j=0}^{mrow2-1} min1[c+i,\ r+j] \bullet min2[(mcol2-1)-i,\ (mrow2-1)-j]$$

for c = 0 to `mcol1+mcol2-1` and r = 0 to `mrow2-1`

Domain

–1.0 to +1.0

conv2d3x3

2-D convolution with 3 x 3 matrix

Synopsis

```
#include <filter.h>
void conv2d3x3_fr16(min1, mrow1, mcol1, min2, mout )
const fract16 *min1;        /* pointer to input matrix 1     */
int mrow1;                  /* number of rows in matrix 1    */
int mcol1;                  /* number of columns in matrix 1 */
const fract16 *min2;        /* pointer to input matrix 2     */
fract16 *mout;              /* pointer to output matrix      */
```

Description

The `conv2d3x3` function computes the two-dimensional circular convolution of matrix `min1` (size `mrow1` x `mcol1`) with matrix `min2` (size 3 x 3).

Algorithm

Two-dimensional input matrix `min1` is convolved with input matrix `min2`, placing the result in a matrix pointed to by `mout`.

$$mout\,[c, r] = \sum_{i=0}^{2} \sum_{j=0}^{2} min1\,[c + i,\ r + j] \bullet min2\,[2 - i, 2 - j]$$

for c = 0 to `mcol1+2` and r = 0 to `mrow1+2`

Domain

–1.0 to +1.0

copysign

copysign

Synopsis

```
#include <math.h>
float copysignf (float parm1, float parm2)
double copysign (double parm1, double parm2)
long double copysignd (long double parm1, long double parm2)
fract16 copysign_fr16 (fract16 parm1, fract16 parm2)
```

Description

The copysign functions copy the sign of the second argument to the first argument.

Algorithm

```
return( |parm1| * copysignof(parm2))
```

Domain

Full range for type of parameters used.

cot

cotangent

Synopsis

```
#include <math.h>
float cotf (float a)
double cot (double a)
long double cotd (long double a)
```

Description

These functions calculate the cotangent of their argument `a`, which is measured in radians. If `a` is outside of the domain, the functions return 0.

Algorithm

```
c = cot(a)
```

Domain

x = [–9099 ... 9099] for `cotf( )`

x = [-4.21657e8 ... 4.21657e8] for `cotd( )`

countones

count one bits in word

Synopsis

```
#include <math.h>
int countones(int parm)
int lcountones(long parm)
int llcountones(long long int parm)
```

Description

The countones functions count the number of one bits in the argument `parm`.

Algorithm

$$return \quad = \sum_{j=0}^{N-1} bit\ [\ j\ ]\ \ of \quad parm$$

where `N` is the number of bits in `parm`.

crosscoh

cross-coherence

Synopsis

```
#include <stats.h>

void crosscohf (const float a[ ],
                const float b[ ],
                int samples, int lags, float out[ ]);

void crosscoh (const double a[ ],
                const double b[ ],
                int samples, int lags, double out[ ]);

void crosscohd (const long double a[ ],
                const long double b[ ],
                int samples, int lags, long double out[ ]);

void crosscoh_fr16 (const fract16 a[ ],
                const fract16 b[ ],
                int samples, int lags, fract16 out[ ]);
```

Description

The crosscoh functions compute the cross-coherence of two input vectors `a[]` and `b[]`. The cross-coherence is the cross-correlation minus the product of the mean of `a` and the mean of `b`. The length of the input vectors is given by `samples`. The functions return the result in the array `out` with `lags` elements.

Algorithm

$$c_k = \frac{1}{n} * \left(\sum_{j=0}^{n-k-1} (a_j - \overline{a}) * (b_{j+k} - \overline{b}) \right)$$

where k = {0,1,...,m-1}, $\overline{a}$ is the mean value of input vector `a`, and $\overline{b}$ is the mean value of input vector `b`.

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `crosscohf ( )`
-1.7×10^{308} to $+1.7 \times 10^{308}$	for `crosscohd ( )`
–1.0 to +1.0	for `crosscoh_fr16 ( )`

crosscorr

cross-correlation

Synopsis

```
#include <stats.h>

void crosscorrf (const float a[ ],
                 const float b[ ],
                 int samples, int lags, float out[ ]);

void crosscorr (const double a[ ],
                const double b[ ],
                int samples, int lags, double out[ ]);

void crosscorrd (const long double a[ ],
                 const long double b[ ],
                 int samples, int lags, long double out[ ]);

void crosscorr_fr16 (const fract16 a[ ],
                 const fract16 b[ ],
                 int samples, int lags, fract16 out[ ]);
```

Description

The crosscorr functions perform a cross-correlation between two signals. The cross-correlation is the sum of the scalar products of the signals in which the signals are displaced in time with respect to one another. The signals to be correlated are given by the input vectors `a[]` and `b[]`. The length of the input vectors is given by `samples`. The functions return the result in the array `out` with `lags` elements.

Cross-correlation is used in signal processing applications such as speech analysis.

Algorithm

$$c_k = \frac{1}{n} * \left(\sum_{j=0}^{n-k-1} a_j * b_{j+k} \right)$$

where k = {0,1,...,m-1}

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `crosscorrf ( )`
-1.7×10^{308} to $+1.7 \times 10^{308}$	for `crosscorrd ( )`
–1.0 to +1.0	for `crosscorr_fr16 ( )`

csub

complex subtraction

Synopsis

```
#include <complex.h>
complex_float csubf (complex_float a, complex_float b);
complex_double csub (complex_double a, complex_double b);
complex_long_double csubd (complex_long_double a,
                           complex_long_double b);
complex_fract16 csub_fr16 (complex_fract16 a, complex_fract16 b);
```

Description

The csub functions compute the complex subtraction of two complex inputs, `a` and `b`, and return the result.

Algorithm

```
Re(c) = Re(a) – Re(b)
Im(c) = Im(a) – Im(b)
```

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `csubf ( )`
-1.7×10^{308} to $+1.7 \times 10^{308}$	for `csubd ( ), csub ( )`
–1.0 to 1.0	for `csub_fr16 ( )`

fir

finite impulse response filter

Synopsis

```
#include <filter.h>

void fir_fr16(x,y,n,s)
const fract16 x[];      /* Input sample vector x               */
fract16 y[];            /* Output sample vector y              */
int n;                  /* Number of input samples             */
fir_state_fr16 *s;      /* Pointer to filter state structure   */
```

The function uses the following structure to maintain the state of the filter.

```
typedef struct
{
   fract16 *h,          /* filter coefficients              */
   fract16 *d,          /* start of delay line              */
   fract16 *p,          /* read/write pointer               */
   int k;               /* number of coefficients           */
   int l;               /* interpolation/decimation index   */
} fir_state_fr16;
```

Description

The `fir_fr16` function implements a finite impulse response (FIR) filter. The function generates the filtered response of the input data `x` and stores the result in the output vector `y`. The number of input samples and the length of the output vector are specified by the argument `n`.

The function maintains the filter state in the structured variable `s` which must be declared and initialized before calling the function. The macro `fir_init`, in the `filter.h` header file, is available to initialize the structure.

It is defined as:

```
#define fir_init(state, coeffs, delay, ncoeffs. index) \
   (state).h = (coeffs);   \
   (state).d = (delay);    \
   (state).p = (delay);    \
   (state).k = (ncoeffs);  \
   (state).l = (index)
```

The characteristics of the filter (passband, stopband, and so on) are dependent upon the number of filter coefficients and their values. A pointer to the coefficients should be stored in `s->h`, and `s->k` should be set to the number of coefficients.

Each filter should have its own delay line which is a vector of type `fract16` and whose length is equal to the number of coefficients. The vector should be initially cleared to zero and should not otherwise be modified by the user program. The structure member `s->d` should be set to the start of the delay line, and the function uses `s->p` to keep track of its current position within the vector.

The structure member `s->l` is not used by `fir_fr16`. This field is normally set to an interpolation/decimation index before calling either the `fir_interp_fr16` or `fir_decima_fr16` functions.

Algorithm

$$y(i) = \sum_{j=0}^{k-1} h(j) * x(i-j) \text{ for } i = 0,1,..n-1$$

Domain

–1.0 to +1.0

fir_decima

FIR decimation filter

Synopsis

```
#include <filter.h>

void fir_decima_fr16(x,y,n,s)
const fract16 x[];      /* Input sample vector x             */
fract16 y[];            /* Output sample vector y            */
int n;                  /* Number of input samples           */
fir_state_fr16 *s;      /* Pointer to filter state structure */
```

The function uses the following structure to maintain the state of the filter.

```
typedef struct
{
   fract16 *h;          /* filter coefficients             */
   fract16 *d;          /* start of delay line             */
   fract16 *p;          /* read/write pointer              */
   int k;               /* number of coefficients          */
   int l;               /* interpolation/decimation index  */
} fir_state_fr16;
```

Description

The `fir_decima_fr16` function performs an FIR-based decimation filter. It generates the filtered decimated response of the input data `x` and stores the result in the output vector `y`. The number of input samples is specified by the argument `n`, and the size of the output vector should be `n/l` where `l` is the decimation index.

The function maintains the filter state in the structured variable `s`, which must be declared and initialized before calling the function. The macro `fir_init`, in the `filter.h` header file, is available to initialize the structure.

It is defined as:

```
#define fir_init(state, coeffs, delay, ncoeffs, index) \
   (state).h = (coeffs);   \
   (state).d = (delay);    \
   (state).p = (delay);    \
   (state).k = (ncoeffs);  \
   (state).l = (index)
```

The characteristics of the filter are dependent upon the number of filter coefficients and their values, and on the decimation index supplied by the calling program. A pointer to the coefficients should be stored in `s->h`, and `s->k` should be set to the number of coefficients. The decimation index is supplied to the function in `s->l`.

Each filter should have its own delay line which is a vector of type `fract16` and whose length is equal to the number of coefficients. The vector should be initially cleared to zero and should not otherwise be modified by the user program. The structure member `s->d` should be set to the start of the delay line, and the function uses `s->p` to keep track of its current position within the vector.

Algorithm

$$y(k) = \sum_{i=0}^{p-1} x(k * l - i) * h(i)$$

where i = 0,1,...,(n/l) - 1

Domain

–1.0 to + 1.0

fir_interp

FIR interpolation filter

Synopsis

```
#include <filter.h>
void fir_interp_fr16(x,y,n,s)
const fract16 x[];      /* Input sample vector x              */
fract16 y[];            /* Output sample vector y             */
int n;                  /* Number of input samples            */
fir_state_fr16 *s;      /* Pointer to filter state structure  */
```

The function uses the following structure to maintain the state of the filter.

```
typedef struct
{
   fract16 *h;         /* filter coefficients                  */
   fract16 *d;         /* start of delay line                  */
   fract16 *p;         /* read/write pointer                   */
   int k;              /* number of coefficients per polyphase */
   int l;              /* interpolation/decimation index       */
} fir_state_fr16;
```

Description

The `fir_interp_fr16` function performs an FIR-based interpolation filter. It generates the interpolated filtered response of the input data `x` and stores the result in the output vector `y`. The number of input samples is specified by the argument `n`, and the size of the output vector should be `n*l` where `l` is the interpolation index.

The filter characteristics are dependent upon the number of polyphase filter coefficients and their values, and on the interpolation factor supplied by the calling program.

The `fir_interp_fr16` function assumes that the coefficients are stored in the following order:

```
coeffs[(np * ncoeffs) + nc]
```

where: `np = {0, 1, ..., nphases-1}`
`nc = {0, 1, ..., ncoeffs-1}`

A pointer to the coefficients is passed into the `fir_interp_fr16` function via the argument `s`, which is a structured variable that represents the filter state. This structured variable must be declared and initialized before calling the function. The `filter.h` header file contains the macro `fir_init` that can be used to initialize the structure and is defined as:

```
#define fir_init(state, coeffs, delay, ncoeffs, index) \
   (state).h = (coeffs);  \
   (state).d = (delay);   \
   (state).p = (delay);   \
   (state).k = (ncoeffs); \
   (state).l = (index)
```

The interpolation factor is supplied to the function in `s->l`. A pointer to the coefficients should be stored in `s->h`, and `s->k` should be set to the number of coefficients per polyphase filter.

Each filter should have its own delay line which is a vector of type `fract16` and whose length is equal to the number of coefficients in each polyphase. The vector should be cleared to zero before calling the function for the first time and should not otherwise be modified by the user program. The structure member `s->d` should be set to the start of the delay line, and the function uses `s->p` to keep track of its current position within the vector.

Algorithm

$$y_m(k) = \sum_{i=0}^{p/l-1} x(k-i) * h(i*l+m)$$

where i = 0,1,...,n-1

Domain

–1.0 to +1.0

Example

```
#include <filter.h>

#define INTERP_FACTOR 5
#define NSAMPLES 50
#define NCOEFFS 35
#define NPOLY INTERP_FACTOR

/* Coefficients */

fract16 coeffs[NCOEFFS];

/* Input, Output, Delay Line, and Filter State */

fract16 input[NSAMPLES], output[INTERP_FACTOR*NSAMPLES];
fract16 delay[NCOEFFS/NPOLY];

fir_state state;

/* Utility Variables */

int i;

/* Initialize the delay line */

for (i = 0; i < NCOEFFS; i++)
    delay[i] = 0;
```

```
/* Initialize the filter state */
fir_init (state, coeffs, delay, (NCOEFFS/NPOLY), INTERP_FACTOR);
/* Call the fir_interp_fr16 function */
fir_interp_fr16 (input, output, NSAMPLES, &state);
```

gen_bartlett

generate Bartlett window

Synopsis

```
#include <window.h>
void gen_bartlett_fr16(w,a,N)
fract16 w[];   /* Window vector                                   */
int a;         /* Address stride in samples for window vector */
int N;         /* Length of window vector                         */
```

Description

This function generates a vector containing the Bartlett window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`.

The Bartlett window is similar to the Triangle window (see on page 4-107) but has the following different properties:

- The Bartlett window always returns a window with two zeros on either end of the sequence, so that for odd `n`, the center section of an `N+2` Bartlett window equals an `N` Triangle window.
- For even `n`, the Bartlett window is still the convolution of two rectangular sequences. There is no standard definition for the Triangle window for even `n`; the slopes of the Triangle window are slightly steeper than those of the Bartlett window.

Algorithm

$$w[n] = 1 - \left| \frac{n - \frac{N-1}{2}}{\frac{N-1}{2}} \right|$$

where n = {0, 1, 2, ..., N-1}

Domain

a > 0; N > 0

gen_blackman

generate Blackman window

Synopsis

```
#include <window.h>
void gen_blackman_fr16(w,a,N)
fract16 w[];   /* Window vector                                    */
int a;         /* Address stride in samples for window vector */
int N;         /* Length of window vector                          */
```

Description

This function generates a vector containing the Blackman window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`.

Algorithm

$$w[n] = 0.42 - 0.5\cos\left(\frac{2\pi n}{N-1}\right) + 0.08\cos\left(\frac{4\pi n}{N-1}\right)$$

where n = {0, 1, 2, ..., N-1}

Domain

a > 0; N > 0

gen_gaussian

generate Gaussian window

Synopsis

```
#include <window.h>
void gen_gaussian_fr16(w,alpha,a,N)
fract16 w[];   /* Window vector                                  */
float alpha;   /* Gaussian alpha parameter                       */
int a;         /* Address stride in samples for window vector    */
int N;         /* Length of window vector                        */
```

Description

This function generates a vector containing the Gaussian window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`.

Algorithm

$$w(n) = \exp\left[-\frac{1}{2}\left(\alpha\frac{n - N/2 - 1/2}{N/2}\right)^2\right]$$

where n = {0, 1, 2, ..., N-1} and α is an input parameter.

Domain

a > 0; N > 0; α > 0.0

gen_hamming

generate Hamming window

Synopsis

```
#include <window.h>
void gen_hamming_fr16(w,a,N)
fract16 w[];   /* Window vector                                    */
int a;         /* Address stride in samples for window vector */
int N;         /* Length of window vector                          */
```

Description

This function generates a vector containing the Hamming window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`.

Algorithm

$$w[n] = 0.54 - 0.46\cos\left(\frac{2\pi n}{N-1}\right)$$

where n = {0, 1, 2, ..., N-1}

Domain

a > 0; N > 0

gen_hanning

generate Hanning window

Synopsis

```
#include <window.h>
void gen_hanning_fr16(w,a,N)
fract16 w[];   /* Window vector                                  */
int a;         /* Address stride in samples for window vector */
int N;         /* Length of window vector                        */
```

Description

This function generates a vector containing the Hanning window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`. This window is also known as the Cosine window.

Algorithm

$$w[n] = 0.5 - 0.5\cos\left(\frac{2\pi n}{N-1}\right)$$

where n = {0, 1, 2, ..., N-1}

Domain

a > 0; N > 0

gen_harris

generate Harris window

Synopsis

```
#include <window.h>
void gen_harris_fr16(w,a,N)
fract16 w[];   /* Window vector                                    */
int a;         /* Address stride in samples for window vector */
int N;         /* Length of window vector                          */
```

Description

This function generates a vector containing the Harris window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`. This window is also known as the Blackman-Harris window.

Algorithm

$$w[n] = 0.35875 - 0.48829 * \cos\left(\frac{2\pi n}{N-1}\right) + 0.14128 * \cos\left(\frac{4\pi n}{N-1}\right) - 0.01168 * \cos\left(\frac{6\pi n}{N-1}\right)$$

where n = {0, 1, 2, ..., N-1}

Domain

a > 0; N > 0

gen_kaiser

generate Kaiser window

Synopsis

```
#include <window.h>
void gen_kaiser_fr16(w,beta,a,N)
fract16 w[];   /* Window vector                                  */
float beta;    /* Kaiser beta parameter                          */
int a;         /* Address stride in samples for window vector */
int N;         /* Length of window vector                        */
```

Description

This function generates a vector containing the Kaiser window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`. The β value is specified by parameter `beta`.

Algorithm

$$w[n] = \frac{I_0\left[\beta\left(1-\left[\frac{n-\alpha}{\alpha}\right]^2\right)^{1/2}\right]}{I_0(\beta)}$$

where n = {0, 1, 2, ..., N-1}, α = (N - 1) / 2, and $I_0(\beta)$ represents the zeroth-order modified Bessel function of the first kind.

Domain

$\alpha > 0$; $N > 0$; $\beta > 0.0$

gen_rectangular

generate rectangular window

Synopsis

```
#include <window.h>
void gen_rectangular_fr16(w,a,N)
fract16 w[];   /* Window vector                                  */
int a;         /* Address stride in samples for window vector */
int N;         /* Length of window vector                       */
```

Description

This function generates a vector containing the Rectangular window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`.

Algorithm

$w[n] = 1$

where n = {0, 1, 2, ..., N-1}

Domain

a > 0; N > 0

gen_triangle

generate triangle window

Synopsis

```
#include <window.h>
void gen_triangle_fr16(w,a,N)
fract16 w[];    /* Window vector                                    */
int a;          /* Address stride in samples for window vector */
int N;          /* Length of window vector                          */
```

Description

This function generates a vector containing the Triangle window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`.

Refer to the Bartlett window (on page 4-98) regarding the relationship between it and the Triangle window.

Algorithm

For even `n`, the following equation applies.

$$w[n] = \begin{cases} \dfrac{2n+1}{N} & n < N/2 \\ \dfrac{2N-2n-1}{N} & n > N/2 \end{cases}$$

where n = {0, 1, 2, ..., N-1}

For odd `n`, the following equation applies.

$$w[n] = \begin{cases} \dfrac{2n+2}{N+1} & n < N/2 \\ \dfrac{2N-2n}{N+1} & n > N/2 \end{cases}$$

where n = {0, 1, 2, ..., N-1}

Domain

$a > 0; N > 0$

gen_vonhann

generate Von Hann window

Synopsis

```
#include <window.h>
void gen_vonhann_fr16(w,a,N)
fract16 w[];   /* Window vector                                  */
int a;         /* Address stride in samples for window vector */
int N;         /* Length of window vector                       */
```

Description

This function is identical to the Hanning window (see on page 4-103).

Domain

a > 0; N > 0

histogram

histogram

Synopsis

```
#include <stats.h>

void histogramf (const float in[ ], int out[ ],
                 float max,
                 float min, int nsamples, int nbins);

void histogram (const double in[ ], int out[ ],
                double max,
                double min, int nsamples, int nbins);

void histogramd (const long double in[ ], int out[ ],
                 long double max,
                 long double min, int nsamples, int nbins);

void histogram_fr16 (const fract16 in[ ], int out[ ],
                 fract16 max,
                 fract16 min, int nsamples, int nbins);
```

Description

The histogram functions compute a histogram of the input vector `in[ ]` that contains `nsamples` samples, and store the result in the output vector `out`.

The minimum and maximum value of any input sample is specified by `min` and `max`, respectively. These values are used by the function to calculate the size of each bin as `(max - min) / nbins`, where `nbins` is the size of the output vector `out`.

Any input value that is outside the range `[ min, max )` exceeds the boundaries of the output vector and is discarded.

To preserve maximum performance while performing out-of-bounds checking, the `histogram_fr16` function allocates a temporary work area on the stack. The work area is allocated with `(nbins + 2)` elements and the stack may therefore overflow if the number of bins is sufficiently large. The size of the stack may be adjusted by making appropriate changes to the `.LDF` file.

Algorithm

Each input value is adjusted by `min`, multiplied by `1/binsize`, and rounded. The appropriate bin in the output vector is then incremented.

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `histogramf ( )`
-1.7×10^{308} to $+1.7 \times 10^{308}$	for `histogramd ( )`
−1.0 to +1.0	for `histogram_fr16 ( )`

ifft

N-point radix-2 inverse FFT

Synopsis

```
#include <filter.h>
void ifft_fr16(in[], t[], out[], w[], wst, n, block_exponent,
               scale_method)
const complex_fract16 in[];   /* input sequence                */
complex_fract16 t[];          /* temporary working buffer      */
complex_fract16 out[];        /* output sequence               */
const complex_fract16 w[];    /* twiddle sequence              */
int wst;                      /* twiddle factor stride         */
int n;                        /* number of FFT points          */
int block_exponent;           /* block exponent of output data */
int scale_method;             /* scaling method desired:
                                 0-none, 1-static, 2-dynamic   */
```

Description

This function transforms the frequency domain complex input signal sequence to the time domain by using the radix-2 Fast Fourier Transform.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n`, where `n` represents the number of points in the FFT. To avoid potential data bank collisions the input and temporary buffers should be allocated in different memory banks; this results in improved run-time performance. If the input data can be overwritten, the optimum memory usage can be achieved by also specifying the input array as the output array.

The twiddle table is passed in the argument `w`, which must contain at least `n/2` twiddle coefficients. The function `twidfftrad2_fr16` may be used to initialize the array. If the twiddle table contains more coefficients than needed for a particular call on `ifft_fr16`, then the stride factor has to be set appropriately; otherwise it should be set to 1.

The arguments `block_exponent` and `scale_method` have been added for future expansion. These arguments are ignored by the function. To avoid overflow the function scales the output by `1/n`.

Algorithm

$$x(n) = \frac{1}{N}\sum_{k=0}^{N-1} X(k)W_N^{-nk}$$

The implementation uses core FFT functions. To get the inverse effect, the function first swaps the real and imaginary parts of the input, performs the direct radix-2 transformation, and finally swaps the real and imaginary parts of the output.

Domain

Input sequence length `n` must be a power of 2 and at least 16.

ifftrad4

N-point radix-4 inverse input FFT

Synopsis

```
#include <filter.h>
void ifftrad4_fr16 (in[], t[], out[], w[], wst, n,
                    block_exponent, scale_method)
const complex_fract16 in[];   /* input sequence                */
complex_fract16 t[];          /* temporary working buffer      */
complex_fract16 out[];        /* output sequence               */
const complex_fract16 w[];    /* twiddle sequence              */
int wst;                      /* twiddle factor stride         */
int n;                        /* number of FFT points          */
int block_exponent;           /* block exponent of output data */
int scale_method;             /* scaling method desired:
                                 0-none, 1-static, 2-dynamic   */
```

Description

This function transforms the frequency domain complex input signal sequence to the time domain by using the radix-4 Inverse Fast Fourier Transform.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n`, where `n` represents the number of points in the FFT. Memory bank collisions, which have an adverse effect on run-time performance, may be avoided by allocating all input and working buffers to different memory banks. If the input data can be overwritten, the optimum memory usage can be achieved by also specifying the input array as the output array.

The twiddle table is passed in the argument `w`, which must contain at least `3/4n` twiddle factors. The function `twidfftrad4_fr16` may be used to initialize the array. If the twiddle table contains more factors than needed for a particular call on `ifftrad4_fr16`, then the stride factor has to be set appropriately; otherwise it should be set to 1.

The arguments `block_exponent` and `scale_method` have been added for future expansion. These arguments are ignored by the function. To avoid overflow, the function performs static scaling by first dividing the input by `n`.

Algorithm

$$x(n) = \frac{1}{N}\sum_{k=0}^{N-1} X(k)W_N^{-nk}$$

The implementation uses core FFT functions. To get the inverse effect, the function first swaps the real and imaginary parts of the input, performs the direct radix-4 transformation, and finally swaps the real and imaginary parts of the output.

Domain

Input sequence length `n` must be a power of 4 and at least 16.

ifft2d

n x n point 2-D inverse input FFT

Synopsis

```
#include <filter.h>
void ifft2d_fr16(*in, *t, *out, w[], wst, n, block_exponent,
                 scale_method)
const complex_float *in;     /* pointer to input matrix a[n][n]   */
complex_fract16 *t;          /* pointer to working buffer t[n][n] */
const complex_fract16 w[];   /* twiddle sequence                  */
int wst;                     /* twiddle factor stride             */
int n;                       /* number of FFT points              */
int block_exponent;          /* block exponent of output data     */
int scale_method;            /* scaling method desired:
                                0-none, 1-static, 2-dynamic       */
```

Description

This function computes a two-dimensional Inverse Fast Fourier Transform of the complex input matrix `a[n][n]` and stores the result to the complex output matrix `c[n][n]`.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n*n`, where `n` represents the number of points in the FFT. Memory bank collisions, which have an adverse effect on run-time performance, may be avoided by allocating all input and working buffers to different memory banks. If the input data can be overwritten, the optimum memory usage can be achieved by also specifying the input array as the output array.

The twiddle table is passed in the argument `w`, which must contain at least `n` twiddle factors. The function `twidfft2d_fr16` may be used to initialize the array. If the twiddle table contains more factors than needed for a particular call on `ifft2d_fr16`, then the stride factor has to be set appropriately; otherwise it should be set to 1.

The arguments `block_exponent` and `scale_method` have been added for future expansion. These arguments are ignored by the function. To avoid overflow the function performs static scaling by dividing the input by `n*n`.

Algorithm

$$c(i,j) = \frac{1}{n^2}\sum_{k=0}^{n-1}\sum_{l=0}^{n-1} a(k,l) * e^{2\pi j(i*k+j*l)/n}$$

where i={0,1,...,n-1}, j={0,1,2,...,n-1}

Domain

Input sequence length `n` must be a power of 2 and at least 16.

iir

infinite impulse response filter

Synopsis

```
#include <filter.h>
void iir_fr16 (x,y,n,s)
const fract16 x[];   /* Input sample vector x             */
fract16 y[];         /* Output sample vector y            */
int n;               /* Number of input samples           */
iir_state_fr16 *s;   /* Pointer to filter state structure */
```

The function uses the following structure to maintain the state of the filter.

```
typedef struct
{
   fract16 *c;         /* coefficients             */
   fract16 *d;         /* start of delay line      */
   int k;              /* number of biquad stages  */
} iir_state_fr16;
```

Description

The `iir_fr16` function implements a biquad, canonical form, infinite impulse response (IIR) filter. It generates the filtered response of the input data `x` and stores the result in the output vector `y`. The number of input samples and the length of the output vector are specified by the argument `n`.

The function maintains the filter state in the structured variable `s`, which must be declared and initialized before calling the function. The macro `iir_init`, in the `filter.h` header file, is available to initialize the structure and is defined as:

```
#define iir_init(state, coeffs, delay, stages) \
   (state).c = (coeffs);    \
   (state).d = (delay);     \
   (state).k = (stages)
```

The characteristics of the filter are dependent upon filter coefficients and the number of stages. Each stage has five coefficients which must be stored in the order `A2`, `A1`, `B2`, `B1`, and `B0`. The value of `A0` is implied to be `1.0` and `A1` and `A2` should be scaled accordingly. This requires that the value of the `A0` coefficient be greater than both `A1` and `A2` for all the stages. The function `iirdf1_fr16` (see on page 4-121) implements a direct form I filter, and does not impose this requirement; however, it does assume that the `A0` coefficients are 1.0.

A pointer to the coefficients should be stored in `s->c`, and `s->k` should be set to the number of stages.

The `iir_fr16` function is implemented using a direct form II algorithm. When importing coefficients from a filter design tool that employs a transposed direct form II, the `A1` and `A2` coefficients have to be negated. For example, if a filter design tool returns `A = [1.0, 0.2, -0.9]`, then the A coefficients have to be modified to `A = [1.0, -0.2, 0.9]`.

Each filter should have its own delay line which is a vector of type `fract16` and whose length is equal to twice the number of stages. The vector should be initially cleared to zero and should not otherwise be modified by the user program. The structure member `s->d` should be set to the start of the delay line.

Algorithm

$$H(z) = \frac{B_0 + B_1 z^{-1} + B_2 z^{-2}}{1 - A_1 z^{-1} - A_2 z^{-2}}$$

where

$$D_m = A_2 * D_{m-2} + A_1 * D_{m-1} + x_m$$
$$Y_m = B_2 * D_{m-2} + B_1 * D_{m-1} + B_0 * D_m$$

where m = {0, 1, 2, ..., n-1}

Domain

–1.0 to +1.0

iirdf1

direct form I impulse response filter

Synopsis

```
#include <filter.h>
void iirdf1_fr16 (const fract16 x[ ],
                  fract16 y[ ],
                  int n,
                  iirdf1_fr16_state *s);
```

The function uses the following structure to maintain the state of the filter.

```
typedef struct
{
   fract16 *c;     /* coefficients                    */
   fract16 *d;     /* start of delay line             */
   fract16 *p;     /* read/write pointer              */
   int k;          /* 2*number of biquad stages + 1   */
} iirdf1_fr16_state;
```

Description

The `iirdf1_fr16` function implements a biquad, direct form I, infinite impulse response (IIR) filter. It generates the filtered response of the input data `x` and stores the result in the output vector `y`. The number of input samples and the length of the output vector is specified by the argument `n`.

The function maintains the filter state in the structured variable `s`, which must be declared and initialized before calling the function. The macro `iirdf1_init`, defined in the `filter.h` header file, is available to initialize the structure.

The macro is defined as:

```
#define iirdf1_init(state, coeffs, delay, stages) \
   (state).c = (coeffs);      \
   (state).d = (delay);       \
   (state).p = (delay);       \
   (state).k = (2*(stages)+1)
```

The characteristics of the filter are dependent upon the filter coefficients and the number of biquad stages. The `A`-coefficients and the `B`-coefficients for each stage are stored in a vector that is addressed by the pointer `s->c`. This vector should be generated by the `coeff_iirdf1_fr16` function (see "coeff_iirdf1" on page 4-74). The variable `s->k` should be set to the expression `(2*stages) + 1`.

Both the `iirdf1_fr16` and `iir_fr16` functions assume that the value of the `A0` coefficients is 1.0, and that all other `A`-coefficients have been scaled according. For the `iir_fr16` function, this also implies that the value of the `A0` coefficient is greater than both the `A1` and `A2` for all stages. This restriction does not apply to the `iirdf1_fr16` function because the coefficients are specified as floating-point values to the `coeff_iirdf1_fr16` function.

Each filter should have its own delay line which is a vector of type `fract16` and whose length is equal to `(4 * stages) + 2`. The vector should be initially cleared to zero and should not otherwise be modified by the user program. The structure member `s->d` should be set to the start of the delay line, and the function uses `s->p` to keep track of its current position within the vector. For optimum performance, coefficient and state arrays should be allocated in separate memory blocks.

The `iirdf1_fr16` function will adjust the output by the scaling factor that was applied to the `A`-coefficients and the `B`-coefficients by the `coeff_iirfd1_fr16` function.

Algorithm

$$H(z) = \frac{B_0 + B_1 z^{-1} + B_2 z^{-2}}{1 - A_1 z^{-1} - A_2 z^{-2}}$$

where:

```
V    = B0 * X(i) + B1 * X(i-1) + B2 * X(i-2)
Y(i) = V - A1 * Y(i-1) - A2 * Y(i-2)
```

where `i = {0, 1, .., n-1}`

Domain

-1.0 to +1.0

Example

```
#include <filter.h>

#define NSAMPLES 50
#define NSTAGES 2

/* Coefficients for the coeff_iirdf1_fr16 function */

const float a_coeffs[(2 * NSTAGES)] = { . . . };
const float b_coeffs[(2 * NSTAGES) + 1] = { . . . };

/* Coefficients for the iirdf1_fr16 function */

fract16 df1_coeffs[(4 * NSTAGES) + 2];

/* Input, Output, Delay Line, and Filter State */

fract16 input[NSAMPLES], output[NSAMPLES];
fract16 delay[(4 * NSTAGES) + 2];
iirdf1_fr16_state state;
int i;
```

```
/* Initialize filter description */

iirdf1_init (state,df1_coeffs,delay,NSTAGES);

/* Initialize the delay line */

for (i = 0; i < ((4 * NSTAGES) + 2); i++)
    delay[i] = 0;

/* Convert coefficients */

coeff_iirdf1_fr16 (a_coeffs,b_coeffs,df1_coeffs,NSTAGES);

/* Call the function */

iirdf1_fr16 (input,output,NSAMPLES,&state);
```

max

maximum

Synopsis

```
#include <math.h>

int max (int parm1, int parm2)
long int lmax (long int parm1, long int parm2)
long long int llmax (long long int parm1, long long int parm2)

float fmaxf (float parm1, float parm2)
double fmax (double parm1, double parm2)
long double fmaxd (long double parm1, long double parm2)

fract16 max_fr16 (fract16 parm1, fract16 parm2)
```

Description

These functions return the larger of their two arguments.

Algorithm

```
if (parm1 > parm2)
   return (parm1)
else
   return(parm2)
```

Domain

Full range for type of parameters.

mean

mean

Synopsis

```
#include <stats.h>

float meanf (const float in[ ], int n);
double mean (const double in[ ], int n);
long double meand (const long double in[ ], int n);

fract16 mean_fr16 (const fract16 in[ ], int n);
```

Description

These functions return the mean of the input array `in[ ]`. The number of elements in the array is `n`.

Algorithm

$$c = \frac{1}{n} * (\sum_{i=0}^{n-1} a_i)$$

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `meanf ( )`
-1.7×10^{308} to $+1.7 \times 10^{308}$	for `meand ( )`
–1.0 to +1.0	for `mean_fr16 ( )`

min

minimum

Synopsis

```
#include <math.h>

int min (int parm1, int parm2)
long int lmin (long int parm1, long int parm2)
long long int llmin (long long int parm1, long long int parm2)

float fminf (float parm1, float parm2)
double fmin (double parm1, double parm2)
long double fmind (long double parm1, long double parm2)

fract16 min_fr16 (fract16 parm1, fract16 parm2)
```

Description

These functions return the smaller of their two arguments.

Algorithm

```
if (parm1 < parm2)
   return( parm1)
else
   return(parm2)
```

Domain

Full range for type of parameters used.

mu_compress

μ-law compression

Synopsis

```
#include <filter.h>
void mu_compress (in, out, n)
const short in[];      /* Input array                          */
short out[];           /* Output array                         */
int n;                 /* Number of elements to be compressed */
```

Description

This function takes a vector of linear 14-bit signed speech samples and performs μ-law compression according to ITU recommendation G.711. Each sample is compressed to 8 bits and is returned in the vector pointed to by `out`.

Algorithm

`C(k)= mu_law compression` of $A(k)$ for k = 0 to n-1

Domain

Content of input array: –8192 to 8191

mu_expand

μ-law expansion

Synopsis

```
#include <filter.h>
void mu_expand (in, out, n)
const short in[];       /* Input array                        */
short out[];            /* Output array                       */
int n;                  /* Number of elements to be expanded */
```

Description

This function inputs a vector of 8-bit compressed speech samples and expands them according to ITU recommendation G.711. Each input value is expanded to a linear 14-bit signed sample in accordance with the μ-law definition and is returned in the vector pointed to `out`.

Algorithm

`C(k)= mu_law expansion` of *A(k)* for k = 0 to n-1

Domain

Content of input array: 0 to 255

norm

normalization

Synopsis

```
#include <complex.h>
complex_float normf (complex_float a)
complex_double norm (complex_double a)
complex_long_double normd (complex_long_double a)
```

Description

These functions normalize the complex input `a` and return the result.

Algorithm

$$\mathrm{Re}(c) = \frac{\mathrm{Re}(a)}{\sqrt{\mathrm{Re}^2(a) + \mathrm{Im}^2(a)}}$$

$$\mathrm{Im}(c) = \frac{\mathrm{Im}(a)}{\sqrt{\mathrm{Re}^2(a) + \mathrm{Im}^2(a)}}$$

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$ for `normf ( )`

-1.7×10^{308} to $+1.7 \times 10^{308}$ for `normd ( )`

polar

construct from polar coordinates

Synopsis

```
#include <complex.h>
complex_float polarf (float magnitude, float phase)
complex_double polar (double magnitude, double phase)
complex_long_double polard (long double magnitude,
                            long double phase)
complex_fract16 polar_fr16 (fract16 magnitude, fract16 phase)
```

Description

These functions transform the polar coordinate, specified by the arguments `magnitude` and `phase`, into a Cartesian coordinate and return the result as a complex number in which the x-axis is represented by the real part, and the y-axis by the imaginary part. The phase argument is interpreted as radians.

For the `polar_fr16` function, the phase must be scaled by 2π and must be in the range `[0x8000, 0x7ff0]`. The value of the phase may be either positive or negative. Positive values are interpreted as an anti-clockwise motion around a circle with a radius equal to the magnitude as shown in Table 4-10.

Table 4-10. Positive Phases

Phase	Radians
0.0	0
0.25(0x2000)	π/2
0.50(0x4000)	π
0.75(0x6000)	3/2π
0.999(0x7ff0)	<2π

Table 4-11 shows how negative values for the phase argument are interpreted as a clockwise movement around a circle.

Table 4-11. Negative Phases

Phase	Radians
-0.25(0xe000)	3/2π
-0.50(0xc000)	π
-0.75(0xa000)	π/2
-1.00(0x8000)	2 π

Algorithm

```
Re(c) = r*cos(θ)
Im(c) = r*sin(θ)
```

where θ is the phase, and `r` is the magnitude

Domain

phase = [–4.3e7 ... 4.3e7] for `polarf( )`
magnitude = -3.4×10^{38} ... $+3.4 \times 10^{38}$ for `polarf( )`

phase = [–8.4331e8 ... 8.4331e8] for `polard( )`
magnitude = -1.7×10^{308} to $+1.7 \times 10^{308}$ for `polard( )`

phase = [–1.0 ...+0.999969] for `polar_fr16( )`
magnitude = [–1.0 ... 1.0) for `polar_fr16( )`

Example

```
#include <complex.h>

#define PI 3.14159265

complex_fract16 point;
float phase_float;

fract16 phase_fr16;
fract16 mag_fr16;

phase_float = PI;
phase_fr16 = (phase_float / (2*PI)) * 32768.0;
mag_fr16 = 0x0200;

point = polar_fr16 (mag_fr16,phase_fr16);
        /* point.re = 0xfe00 */
        /* point.im = 0x0000 */
```

rfft

N-point radix-2 real input FFT

Synopsis

```
#include <filter.h>
void rfft_fr16(in[], t[], out[], w[], wst, n,
               block_exponent, scale_method)
const fract16 in[];          /* input/output sequence         */
complex_fract16 t[];         /* temporary working buffer      */
complex_fract16 out[];       /* working buffer                */
const complex_fract16 w[];   /* twiddle sequence              */
int wst;                     /* twiddle factor stride         */
int n;                       /* number of FFT points          */
int block_exponent;          /* block exponent of output data */
int scale_method;            /* scaling method desired:
                                0-none, 1-static, 2-dynamic   */
```

Description

This function transforms the time domain real input signal sequence to the frequency domain by using the radix-2 FFT. The function takes advantage of the fact that the imaginary part of the input equals zero, which in turn eliminates half of the multiplications in the butterfly.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n`, where `n` represents the number of points in the FFT. Memory bank collisions, which have an adverse effect on run-time performance, may be avoided by allocating all input and working buffers to different memory banks. If the input data can be overwritten, the optimum memory usage can be achieved by also specifying the input array as the output array, provided that the memory size of the input array is at least `2*n`.

The twiddle table is passed in the argument `w`, which must contain at least `n/2` twiddle factors. The function `twidfftrad2_fr16` may be used to initialize the array. If the twiddle table contains more factors than needed for a particular call on `rfft_fr16`, then the stride factor has to be set appropriately; otherwise it should be set to 1.

The arguments `block_exponent` and `scale_method` have been added for future expansion. These arguments are ignored by the function. To avoid overflow, the function performs static scaling by dividing the input by `1/n`.

Algorithm

See "cfft" on page 4-61 for more information.

Domain

Input sequence length `n` must be a power of 2 and at least 16.

rfftrad4

N-point radix-4 real input FFT

Synopsis

```
#include <filter.h>
void rfftrad4_fr16(in[], t[], out[], w[], wst, n,
                   block_exponent, scale_method)
const fract16 in[];          /* input/output sequence         */
complex_fract16 t[];         /* temporary working buffer      */
complex_fract16 out[];       /* working buffer                */
const complex_fract16 w[];   /* twiddle sequence              */
int wst;                     /* twiddle factor stride         */
int n;                       /* number of FFT points          */
int block_exponent;          /* block exponent of output data */
int scale_method;            /* scaling method desired:
                                0-none, 1-static, 2-dynamic   */
```

Description

This function transforms the time domain real input signal sequence to the frequency domain by using the radix-4 Fast Fourier Transform. The `rfftrad4_fr16` function takes advantage of the fact that the imaginary part of the input equals zero, which in turn eliminates half of the multiplications in the butterfly.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n`, where `n` represents the number of points in the FFT. To avoid potential data bank collisions, the input and temporary buffers should reside in different memory backs; this results in improved run-time performance. If the input data can be overwritten, the optimum memory usage can be achieved by also specifying the input array as the output array, provided that the memory size of the input array is at least `2*n`.

The twiddle table is passed in the argument `w`, which must contain at least `3/4n` twiddle factors. The function `twidfftrad4_fr16` may be used to initialize the array. If the twiddle table contains more factors than needed for a particular call on `rfftrad4_fr16`, then the stride factor has to be set appropriately; otherwise it should be set to 1.

The arguments `block_exponent` and `scale_method` have been added for future expansion. These arguments are ignored by the function. To avoid overflow, the function performs static scaling by dividing the input by `n`.

Algorithm

See "cfftrad4" on page 4-66 for more information.

Domain

Input sequence length `n` must be a power of 4 and at least 16.

rfft2d

n x n point 2-D real input FFT

Synopsis

```
#include <filter.h>
void rfft2d_fr16(*in, *t, *out, w[], wst, n, block_exponent
                 scale_method)
const fract16 *in;            /* pointer to input matrix a[n][n]  */
complex_fract16 *t;           /* pointer to working buffer t[n][n] */
complex_fract16 *out;         /* pointer to output matrix [n][n]  */
const complex_fract16 w[];   /* twiddle sequence                  */
int wst;                      /* twiddle factor stride            */
int n;                        /* number of FFT points             */
int block_exponent;           /* block exponent of output data    */
int scale_method;             /* scaling method desired:
                                 0-none, 1-static, 2-dynamic     */
```

Description

This function computes a two-dimensional Fast Fourier Transform of the real input matrix `a[n][n]`, and stores the result to the complex output matrix `c[n][n]`.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n*n`, where `n` represents the number of points in the FFT. Improved run-time performance can be achieved by allocating the input and temporary arrays in separate memory banks; this avoids any memory bank collisions. If the input data can be overwritten, the optimum memory usage can be achieved by also specifying the input array as the output array, provided that the memory size of the input array is at least `2*n*n`.

The twiddle table is passed in the argument `w`, which must contain at least `n` twiddle coefficients. The function `twidfft2d_fr16` may be used to initialize the array. If the twiddle table contains more coefficients than needed for a particular call on `rfft2d_fr16`, then the stride factor has to be set appropriately; otherwise it should be set to 1.

The arguments `block_exponent` and `scale_method` have been added for future expansion. These arguments are ignored by the function. To avoid overflow, the function scales the output by `n*n`.

Algorithm

$$c(i,j) = \sum_{k=0}^{n-1}\sum_{l=0}^{n-1} a(k,l) * e^{-2\pi j(i*k+j*l)/n}$$

where i={0,1,...,n-1}, j={0,1,2,...,n-1}

Domain

Input sequence length `n` must be a power of 2 and at least 16.

rms

root mean square

Synopsis

```
#include <stats.h>

float rmsf (const float in[ ], int n);
double rms (const double in[ ], int n);
long double rmsd (const long double in[ ], int n);

fract16 rms_fr16 (const fract16 in[ ], int n);
```

Description

These functions return the root mean square of the elements within the input vector `in[ ]`. The number of elements in the vector is `n`.

Algorithm

$$c = \sqrt{\frac{\sum_{i=0}^{n-1} a_i^{\,2}}{n}}$$

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$ for `rmsf ( )`

-1.7×10^{308} to $+1.7 \times 10^{308}$ for `rmsd ( )`

–1.0 to +1.0 for `rms_fr16 ( )`

rsqrt

reciprocal square root

Synopsis

```
#include <math.h>
float rsqrtf (float a)
double rsqrt (double a)
long double rsqrtd (long double a)
```

Description

These functions calculate the reciprocal of the square root of the number `a`. If `a` is negative, the functions return 0.

Algorithm

$c = 1/\sqrt{a}$

Domain

$0.0 \ldots 3.4 \times 10^{38}$ for `rsqrtf ( )`

$0.0 \ldots +1.7 \times 10^{308}$ for `rsqrtd ( )`

twidfftrad2

generate FFT twiddle factors for radix-2 FFT

Synopsis

```
#include <filter.h>
void twidfftrad2_fr16 (complex_fract16 w[], int n)
```

Description

This function calculates complex twiddle coefficients for a radix-2 FFT with `n` points and returns the coefficients in the vector `w`. The vector `w`, known as the twiddle table, is normally calculated once and is then passed to an FFT function as a separate argument. The size of the table must be at least `1/2n`, the number of points in the FFT.

FFTs of different sizes can be accommodated with the same twiddle table. Simply allocate the table at the maximum size. Each FFT has an additional parameter, the "stride" of the twiddle table. To use the whole table, specify a stride of 1. If the FFT uses only half the points of the largest FFT, the stride should be 2 (this takes only every other element).

Algorithm

This function takes FFT length `n` as an input parameter and generates the lookup table of complex twiddle coefficients. The samples are:

$$twid_re(k) = \cos\left(\frac{2\pi}{n}k\right)$$

$$twid_im(k) = -\sin\left(\frac{2\pi}{n}k\right)$$

where k = {0, 1, 2, ..., n/2 - 1}

Domain

The FFT length `n` must be a power of 2 and at least 16.

twidfftrad4

generate FFT twiddle factors for radix-4 FFT

Synopsis

```
#include <filter.h>
void twidfftrad4_fr16 (complex_fract16 w[], int n)
void twidfft_fr16(complex_fract16 w[], int n)
```

Description

The `twidfftrad4_fr16` function initializes a table with complex twiddle factors for a radix-4 FFT. The number of points in the FFT are defined by `n`, and the coefficients are returned in the twiddle table `w`.

The size of the twiddle table must be at least `3/4n`, the length of the FFT input sequence. A table can accommodate several FFTs of different sizes by allocating the table at maximum size, and then using the stride argument of the FFT function to specify the step size through the table.

If the stride is set to 1, the FFT function uses all the table; if your FFT has only a quarter of the number of points of the largest FFT, the stride should be 4.

For efficiency, the twiddle table is normally generated once during program initialization and is then supplied to the FFT routine as a separate argument.

The `twidfft_fr16` routine is provided as an alternative to the `twidfftrad4_fr16` routine and performs the same function.

Algorithm

This function takes FFT length `n` as an input parameter and generates the lookup table of complex twiddle coefficients.

The samples generated are:

$$twid_re(k) = \cos\left(\frac{2\pi}{n}k\right)$$

$$twid_im(k) = \sin\left(\frac{2\pi}{n}k\right)$$

where k = {0, 1, 2, ..., ¾n - 1}

Domain

The FFT length n must be a power of 4 and at least 16.

twidfftf_fr16

generate FFT twiddle factors for a fast FFT

Synopsis

```
#include <filter.h>
void twidfftf_fr16 (complex_float w[ ], int n)
```

Description

The `twidfftf_fr16` function generates complex twiddle factors for the fast radix-4 FFT function `cfftf_fr16` (on page 4-63), and stores the coefficients in the vector `w`. The vector `w`, known as the twiddle table, is normally calculated once and is then passed to the fast FFT as a separate argument. The size of the table must be `3/4n`, the number of points in the FFT.

The same twiddle table may be used to calculate FFTs of different sizes provided that the table is generated for the largest FFT. Each FFT function has a stride parameter that the function uses to stride through the twiddle table. Normally, this stride parameter is set to 1, but to generate a smaller FFT, the argument should be scaled appropriately. For example, if a twiddle table is generated for an FFT with N points, then the same twiddle table may be used to generate a N/4-point FFT, provided that the stride parameter is set to 4, or a N/8-point FFT, if the parameter is set to 8.

The twiddle table generated by the `twidfftf_fr16` function is not compatible with the twiddle table generated by the `twidfftrad4_fr16` function (see on page 4-144).

Algorithm

The function calculates a lookup table of complex twiddle factors. The coefficients generated are:

$$twid_re(k) = \cos\left(\frac{2\pi}{n}k\right)$$

$$twid_im(k) = \sin\left(\frac{2\pi}{n}k\right)$$

where k = {0, 1, 2, ..., 3/4n – 1}

Domain

The number of points in the FFT must be a power of 4 and must be at least 16.

twidfft2d

generate FFT twiddle factors for 2-D FFT

Synopsis

```
#include <filter.h>
void twidfft2d_fr16 (complex_fract16 w[], int n)
```

Description

The `twidfft2d_fr16` function generates complex twiddle factors for a 2-D FFT. The size of the FFT input sequence is given by the argument `n` and the function writes the twiddle factors to the vector `w`, known as the twiddle table.

The size of the twiddle table must be at least `n`, the number of points in the FFT. Normally, the table is only calculated once and is then passed to an FFT function as an argument. A twiddle table may be used to generate several FFTs of different sizes by initializing the table for the largest FFT and then using the stride argument of the FFT function to specify the step size through the table. For example, to generate the largest FFT, the stride is set to 1, and to generate an FFT of half this size the stride is set to 2.

Algorithm

This function takes FFT length `n` as an input parameter and generates the lookup table of complex twiddle coefficients.

The samples generated are:

$$twid_re(k) = \cos\left(\frac{2\pi}{n}k\right)$$

$$twid_im(k) = \sin\left(\frac{2\pi}{n}k\right)$$

where k = {0, 1, 2, ..., n-1}

Domain

The FFT length `n` must be a power of 2 and at least 16.

var

variance

Synopsis

```
#include <stats.h>

float varf (const float in[ ], int n);
double var (const double in[ ], int n);
long double vard (const long double in[ ], int n);

fract16 var_fr16 (const fract16 in[ ], int n);
```

Description

These functions return the variance of the elements within the input vector `in[ ]`. The number of elements in the vector is `n`.

Algorithm

$$c = \frac{n * \sum_{i=0}^{n-1} a_i^2 - (\sum_{i=0}^{n-1} a_i)^2}{n(n-1)}$$

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `varf( )`
-1.7×10^{308} to $+1.7 \times 10^{308}$	for `vard( )`
–1.0 to +1.0	for `var_fr16( )`

zero_cross

count zero crossings

Synopsis

```
#include <stats.h>

int zero_crossf (const float in[ ], int n);
int zero_cross (const double in[ ], int n);
int zero_crossd (const long double in[ ], int n);

int zero_cross_fr16 (fract16 in[ ], int n);
```

Description

The zero_cross functions return the number of times that a signal represented in the input array `in[]` crosses over the zero line. If all the input values are either positive or zero, or they are all either negative or zero, then the functions return a zero.

Algorithm

The actual algorithm is different from the one shown below because the algorithm needs to handle the case where an element of the array is zero. However, this example gives you a basic understanding.

```
if ( a(i) > 0 && a(i+1) < 0 )|| (a(i) < 0 && a(i+1) > 0 )
   the number of zeros is increased by one
```

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `zero_crossf ( )`
-1.7×10^{308} to $+1.7 \times 10^{308}$	for `zero_crossd ( )`
–1.0 to +1.0	for `zero_cross_fr16 ( )`

A ADSP-BF561 PROGRAMMING GUIDE

The Blackfin processor family includes dual-core processors, such as the ADSP-BF561 processor. In addition to other features, dual-core processors add a new dimension to application development. Using a single processor, each processor core can be executing its own program in parallel. The dual-core nature of the processor presents additional challenges to the programmer; this section addresses these challenges within the context of VisualDSP++.

The appendix begins with a brief comparison of the single-core versus dual-core Blackfin processors, before describing VisualDSP++ recommended approaches to application development. Finally, it offers guidelines for developing systems on dual-core Blackfin processors. The appendix expects users to be familiar with multiprocessor development within the VisualDSP++ IDDE and to have an understanding of programming for multiple processors/threads.

All examples given are for the ADSP-BF561 processor.

The appendix contains:

- "Dual-Core Blackfin Architecture Overview" on page A-3
- "Approaches Supported in VisualDSP++" on page A-4
- "Single-Core Application" on page A-5
- "One Application Per Core" on page A-8
- "Single Application/Dual Core" on page A-13
- "Run-Time Library Functions" on page A-18

- "Restrictions On Dual-Core Applications" on page A-20
- "Dual-Core Programming Examples" on page A-21
- "Synchronization Functions" on page A-31

The ADSP-BF561 processor adds a new dimension to application development in the form of a dual-core architecture. Using a single ADSP-BF561 processor, each processor core can be executing its own program in parallel. The dual-core nature of the processor presents additional challenges to the programmer; this document addresses these challenges within the context of VisualDSP++ 4.0.

For the most efficient use of information in this appendix. you should be familiar with the Blackfin architecture, including the ADSP-BF533 processor, and have experience in building and executing C or C++ applications for the Blackfin architecture within the VisualDSP++ environment. The appendix focuses only on the additional considerations necessary for dual-core programming.

Dual-Core Blackfin Architecture Overview

Each dual-core Blackfin processor has two Blackfin cores, A and B, each with its own internal L1 memory. There is a common internal memory shared between the two cores, and both cores share access to external memory.

Each core functions independently: they have their own reset address, Event Vector Table, instruction and data caches, and so on. On reset, core A starts running from its reset address, while core B is disabled. Core B starts running when it is enabled by core A. VisualDSP++ enables Core B when it connects to the EZ-KIT Lite board.

When core B starts running, it starts running its own application, from its own reset address.

The two cores use the `TESTSET` instruction to serialize access to shared resources. The `TESTSET` instruction reads and updates a memory location in an atomic fashion. Applications and libraries can build semaphores and other synchronization mechanisms from this primitive.

Refer to the ADSP-BF561 hardware reference for detailed information on ADSP-BF561 processor's architecture.

Approaches Supported in VisualDSP++

VisualDSP++ supports three different approaches to project developing for dual-core Blackfin processors:

- *Single-core applications*
  In this approach, only core A is used, and core B remains disabled (see "Single-Core Application" on page A-5).
- *One application per core*
  In this approach, each core is treated as a separate processor, built individually. The VisualDSP++ project explicitly builds a `.DXE` file for a particular core. Resource sharing is coarse-grain and is managed by the developer (see "One Application Per Core" on page A-8)
- *One application across both cores*
  In this approach, a hierarchy of VisualDSP++ projects builds a single application that supports both cores. Resource sharing is fine-grain, managed by the linker (see "Single Application/Dual Core" on page A-13)

The following sections describe these approaches in more detail.

The approaches represent increasingly levels of sophistication, with corresponding levels of complexity.

A single-core application allows the processor to be used as a migration path from other Blackfin processors and as a means of running standard and legacy applications with minimal effort. Benchmarks are typical examples. This simplistic approach does not exploit the full potential of the dual-core Blackfin processor but provides the fastest route for getting existing code "up and running."

Having one application per core extends this simplistic approach to use both cores. Effectively, two single-core applications are build independently, and run in parallel on the processor. The shared memory areas,

both internal and external, are each sub-divided into three areas—a section dedicated to core A, a section dedicated to core B, and a shared section. It is left upto the developer to arrange for shared, serialized access to the shared areas from each of the cores.

The single-application/dual-core approach is the most powerful, because it allows for all of the shared memory areas to be used efficiently by both cores. Common code can be placed in shared memory to avoid duplication. Shared data can be placed in shared memory without the need for explicit positioning. This approach allows an expert developer to exercise fine control over the structure of the application, using the VisualDSP++ advanced linker capabilities.

The VisualDSP++ libraries and LDFs provide support for multi-core builds, used by the latter two approaches. This support is available through the `-multicore` compiler switch, and through the `__ADI_MULTICORE` linker macro (see "-multicore" on page 1-40 for more information.

Single-Core Application

The single-core application approach is supported by the "default compiler" Linker Description File (LDF). Whenever the compiler is asked to generate an executable file without specifying an `.LDF` file, the compiler uses a default one for the platform in question. For example,

```
ccblkfn -proc ADSP-BF561 prog.c -o prog.dxe
```

does not specify an `.LDF` file, so the compiler uses the default, whereas:

```
ccblkfn -proc ADSP-BF561 prog.c -o prog.dxe -T ./my.ldf
```

directs the compiler to use `./my.ldf` as the `.LDF` file.

The default compiler .LDF file for the ADSP-BF561 processor is …\VisualDSP\Blackfin\ldf\ADSP-BF561.ldf. It is very similar to the corresponding default LDFs for other Blackfin processors, such as the ADSP-BF533 processor, albeit with a different memory map.

The .LDF file creates a single .DXE file that will run on core A, By default, the same .LDF file is also used for the one-application-per-core build, described in "One Application Per Core" on page A-8.

Shared Memory

The .LDF file divides the 128KB shared L2 internal memory into three areas, where the lower 32 Kbits are reserved for core B, the next 32 Kbits are reserved for core A, and the final 64 Kbits are considered shared between the two cores.

The .LDF file divides the shared L2 internal memory as follows:

- Lowest 32KB: reserved for core B, so not used in this approach;
- Next 32KB: reserved for core A, so usable via section l2_sram_a;
- Most of remaining 64KB: reserved for shared data, so usable via section l2_shared.
- 16 bytes reserved for synchronization locks, so not used by this approach;
- 1KB reserved for second-stage boot loader, not used by LDF/libraries.

Note that much of the internal L2 memory is reserved for either core B or shared use. This is because the same .LDF file is also used for the "one application per core" approach, which is described later. For a single-core application, it may be desirable to customize the .LDF file, so that all of L2 internal memory is available for core A, although this will complicate migration towards a multi-core solution.

To place code or data into the area reserved for core A, place them into the `l2_sram_a` section.

External memory is partitioned between the cores. SDRAM Bank 0 is allocated to core A, and can be used via the section `sdram_data`. SDRAM Bank 1 is allocated to core B, so is not used by this approach.

Synchronization

Synchronization is not necessary for the single-core approach. The `.LDF` file still reserves a section of internal L2 memory for synchronization locks, because a lock is necessary for communicating with the emulator during I/O operations.

Cache, Startup and Events

For a single-core application, normal cache configuration and event handling is used, such as for the ADSP-BF533 processor. The only difference is that the `.LDF` file maps into a Cache Protection Lookaside Buffers (CPLBs) configuration table explicitly for each core. Where `ADSP-BF533.ldf` links against `cplbtab533.doj`, a single-core ADSP-BF561 processor application would link against `cplbtab561a.doj`, for core A's CPLB configuration.

The run-time header executed on startup is a generic routine that has been assembled for the ADSP-BF561 processor. It behaves in the same manner as for other Blackfin platforms, except that it makes no attempt to modify the clock speed. It enables cache and interrupts and exceptions in the same fashion as for other Blackfin processors.

Creating Customized LDFs

When creating a customized `.LDF` file for a single-core application, do not use the Expert Linker (EL) to do the initial creation because EL always creates customized `.LDF` files based on the single-application/dual-core

approach. Instead, to create a customized LDF for a single-core application, copy the default LDF for your platform (for example, the `ADSP-BF561.ldf` file from the `...\VisualDSP\Blackfin\ldf` directory into your project directory), "add" it, and then open it for editing (either in the Expert Linker or in a source window, as desired).

One Application Per Core

Like the single-core application approach, the one-application-per-core ("per-core" hereafter, for simplicity) approach uses the default compiler `.LDF` file. This `.LDF` file builds one application for each invocation, either for core A or for core B, according to command-line options. To produce the two applications, first build the application for one core, and then build the application for the second core.

For example,

```
ccblkfn -proc ADSP-BF561 -flags-link -MDCOREA -o p0.dxe a1.c a2.c
ccblkfn -proc ADSP-BF561 -flags-link -MDCOREB -o p1.dxe b1.c b2.c
```

which would build two applications— `p0.dxe` for core A and `p1.dxe` for core B. Note the use of the `-multicore` compiler switch, which indicates to the compiler and the `.LDF` file that this build should use the library variants that have been built with the necessary multi-core locking enabled.

The `COREA` and `COREB` linker flags define preprocessor macros that select alternative `PROCESSOR` directives in the `.LDF` file. If neither `COREA` nor `COREB` is defined, the `.LDF` file automatically defines `COREA`, and links for core A. This is how the single-core application (described in "Single-Core Application" on page A-5) is implemented.

Shared Memory

The memory map for the .LDF file defines all of the internal memories for both cores, although the PROCESSOR section only uses the areas that are defined for the currently-selected core. Thus, while COREA is defined, the .LDF file maps the "L2_sram_a" section into the middle 32 Kbit section of the L2 internal memory, When COREB is defined, it maps the "L2_sram_b" section into the lowest 32 Kbits of the L2 internal memory. In this manner, the two separate builds can map code or data into the common L2 internal memory without conflicting.

External memory is partitioned between the cores. SDRAM Bank 0 is allocated to core A, and can be used via the "sdram_data" section. Most of SDRAM Bank 1 is allocated to core B, and can similarly be used through the "sdram_data" section. Part of Bank 1 is mapped to both cores as a shared resource, and can be used via the "sdram_shared" section.

Sharing Data

The default .LDF file provides two shared data areas, "l2_shared" and "sdram_shared", which are in L2 Memory and SDRAM, respectively. For the per-core approach, there are a number of methods by which data can be shared within this common area. Some are suggested here; others may be appropriate.

- Use the RESOLVE command:
  The LDF syntax allows a symbol to be explicitly resolved. The example LDFs use RESOLVE to map the symbol "start" to the Reset address, giving the literal address explicitly.

- RESOLVE can also be given a .DXE filename instead of an address; the linker will examine the .DXE file to determine to where the symbol was resolved in that .DXE file , and will use the same address for the symbol in the new output file, if a local definition does not exist. Therefore, you can add RESOLVE directives to the PROCESSOR direc-

tive for core B that specify core A's .DXE file. When building core B's .DXE file, the linker will resolve the symbols to the same address used for core A if a local definition does not exist in core B composite objects. This is the recommended approach, and is demonstrated in the examples that follow.

- Use a global structure:
  The structure would contain all of the shared data, with the definition common to both applications via a shared header file. Map the structure into the shared area within the .LDF file. As the structure is the only (or first) item mapped into the section, it is located at the start of the section.

- If your application is using the VisualDSP++ Kernel (VDK), you can use the VDK message-passing facilities to pass the address of common data from one core to the other.

In all cases, data shared between the two applications must be declared as volatile, so that the compiler does not cache values in registers during times when the other core might be updating the value.

Note that when data with initial values are mapped into the shared space, each application—each .DXE file—will contain a copy of that data.

The data caches within cores A and B do not maintain coherence, so two alternatives are available:

- Do not enable data caching for shared areas

- After finishing an access, but before releasing the data to be used by the other core, flush the data from the cache and invalidate the corresponding cache entries

Sharing Code

To share code between applications:

- Map the functions to the shared area, in the application for core A
- In the application for core B, declare the functions as external (via `extern`)
- Add `RESOLVE` directives to the `PROCESSOR` directive for core B, giving the functions' external names and the filename for core A `.DXE` file

Shared Code With Private Data

It is sometimes desirable for a function to maintain its own private data. In a single-threaded, single-core application, this can be achieved merely by declaring the data as “static”.

In a dual-core application where the code is shared, this would mean the same data would be used by both cores. If the intention is that each core has its own instance of the private data, you can use library routines provided with VisualDSP++ to allocate private copies of the data. These routines are described in detail in “adi_obtain_mc_slot” on page 3-60.

Synchronization

Synchronization functions exist in the run-time library for claiming and releasing a lock variable. They are described in more detail in “adi_acquire_lock” on page 3-57.

```
#include <ccblkfn.h>
void adi_acquire_lock(testset_t *t);
int adi_try_lock(testset_t *t);
void adi_release_lock(testset_t *t);
int adi_core_id(void);
```

The VisualDSP++ Kernel (VDK) is an example of the per-core approach. In VDK, each core runs a separate executable, and each executable runs multiple threads. VDK also supports messaging between the two cores which provides an alternative mechanism for sharing data between cores. VDK does not offer any additional mechanisms for sharing code between the two applications.

Cache, Startup and Events

Each core has its own caches and its own cache configuration table. These are linked in by the `.LDF` file according to whether `COREA` or `COREB` is defined. `COREA` links against `cplbtab561a.doj`, while `COREB` links against `cplbtab561b.doj`.

Each application has its own copy of the `___cplb_ctrl` cache configuration variable. Each application also has its own definitions of the "guard" symbols that the `.LDF` file defines to indicate whether L1 SRAM spaces are available for cache use, so the two applications can run with entirely independent cache configurations. The section "`cplb_code`" must be mapped into L1 Instruction memory so that the CPLB configuration routines can access these core-specific guard symbols.

The startup code is the same for the two cores, however. That is, each application in the per-core approach receives its own copy of the same startup code, resolved to the `Reset` address of that core. In particular, the default startup code does not include any functionality to allow core A to enable core B; the EZ-KIT Lite systems and emulator arrange for core B to be running when its application is loaded.

Each core registers its own event handler (for CPLB events, if requested), and handles interrupts and exceptions separately. The two applications can have separate event masks. Signals can be passed between the two applications by triggering interrupts via the System Interrupt Controller.

The run-time library allows interrupt handlers to be registered, but does not provide any direct support for the System Interrupt Controller, or for raising events at that level.

Creating Customized LDFs

Creating a customized `.LDF` file for the per-core approach is done in the same fashion as for the single-core approach. Namely, explicitly copy the `ADSP-BF561.ldf` file into the project, "add" it, and then start editing it. When editing the local copy, it is essential to be aware that there are two distinct `PROCESSOR` directives, dependent on the definitions of `COREA` and `COREB`.

Single Application/Dual Core

This approach generates a single application with just one build process. The application is divided into three components: the two individual cores and the shared memory (all common memory is treated as one, for the purposes of the build).

The single application/dual core approach ("single/dual" hereafter) allows a more complex application to be built, since the three major components are produced during a single linking process that resolves all symbols at once. This allows code and data in the shared memories to be referenced directly from the cores, allowing the cores to use the same instance of a function or data item.

This sharing process makes use of more advanced linker facilities that are not normally required or employed for single applications that run on a single core. These extra capabilities can present a steep learning curve for those new to cross-system linking. Therefore, the single/dual approach adopts a set of conventions to assist in the development of dual-core appli-

cations. The .LDF files provided with VisualDSP++ rely on these conventions, for simplicity. The advanced developer may choose alternative approaches by using entirely customized .LDF files.

Project Conventions

The conventions are as follows.

- The application is arranged as a hierarchy of VisualDSP++ projects, with the final application being the top-level project. This top-level project is of type “DSP executable”.
- Beneath the top-level project, there are four sub-projects, for core A, core B, shared internal L2 memory, and shared external memory areas, respectively. These sub-projects are of type “DSP library.”
- The sub-projects create individual files called corea.dlb, coreb.dlb, sml2.dlb and sml3.dlb, respectively.
- The top-level project links against the libraries generated by the sub-projects, resolving symbols across all of the system at once, and produces three output files—p0.dxe, p1.dxe and sml2.sm. These files may be loaded into the Blackfin processor.

By dividing the application into individual libraries, it is simpler to arrange for a given part of the application to reside within a particular core or within a particular shared memory.

Establishing a convention for filenames (p0.dxe, sml2.dlb, and so on) means that the .LDF file in the top-level project can be utilize the output of a sub-project without needing customization.

Multi-Core Linking

The single/dual approach makes use of advanced linker facilities to resolve cross-references between the cores and shared memories. Each core is described by a PROCESSOR directive, and the two shared memory areas (the

internal L2 memory and the external memory) are described by a single SHARED_MEMORY directive. The PROCESSORs use LINK_AGAINST directives to indicate that they must use the SHARED_MEMORY region to resolve external references. However, there is a defined order in which the linker will resolve symbols:

1. The linker will first use any local definition to resolve any local reference.
2. If a local definition cannot be found, the linker will search the list of libraries, in order, and use the first definition found. The LDFs for the single/dual approach arrange for the linker to first search the library defined by the sub-project, before searching the standard run-time libraries.
3. Finally, if no definition can be found in the libraries, the linker will search in other output files as specified in LINK_AGAINST directives, in order. If one of those output files contains a definition for the symbol, that definition will be used.

Given these rules, and that the .LDF file causes the cores PROCESSORs to LINK_AGAINST the SHARED_MEMORY, this effectively means that the shared memory area acts as another library in the list, albeit one that is always the last to be searched.

Creating the LDF

The single/dual approach requires a custom .LDF file. This is because the default .LDF files for the dual-core Blackfin processors are designed for the simpler single-core and per-core approaches. The Expert Linker templates, however, are all designed for the single/dual approach. Use the Expert Linker to create an .LDF file for your top-level project. It is not necessary to modify the .LDF file in any way, once created. The sub-projects do not require .LDF files.

Shared Memory

Code and data can be mapped into internal L2 memory by placing them into the "sml2" sub-project. The .LDF file links the SHARED_MEMORY area against the library produced by this sub-project. The usual sections (program, data1, constdata, and so on) are mapped, as is "l2_sram".

Code and data can be mapped into the external memory by placing them into the "sml3" sub-project.

Shared Data

Data items that are to be accessed from both cores can be shared by placing the definitions of the data into the one of the "sml2" and "sml3" sub-projects, and declaring the data as "extern" in the corea.dlb and coreb.dlb projects. When the linker creates p0.dxe and p1.dxe, the external references will be resolved by the definition in sub-project's library.

Such shared items must be declared as "volatile" to ensure that the compiler does not retain copies of the data in registers while the other core may be updating them.

Sharing Code in L2 Memory

Code to be shared in L2 internal memory must be placed into the sub-project that creates the sml2.dlb library. The sub-projects for corea.dlb and coreb.dlb can call the functions in sml2.dlb. The linker resolves the references during build.

Because the linker resolves references via libraries (including the run-time libraries) before using the LINK_AGAINST directives, if a library function is used directly by one of the cores and is also used by shared memory, then that function will be duplicated: the core will receive a local copy of the

function because a definition is found in the library before a search is attempted in the shared memory. The shared memory will also receive a local copy for the same reason.

If run-time library functions are to be shared, then they must only be directly referenced via a function that is defined in one of the `sml2` or `sml3` sub-projects.

Synchronization

Synchronization between the cores can be achieved as for the per-core approach, using "adi_acquire_lock" on page 3-57. The synchronization lock variables must be defined in the `sml2.dlb` or `sml3.dlb` sub-projects, so that it is mapped into the shared memory.

Cache, Startup and Events

The `.LDF` files for the single/dual approach map a copy of the startup code into each core, resolving the copies to the `Reset` addresses of the cores. Startup, cache configuration, and events are as for the per-core approach.

The `.LDF` files also map the "`cplb_code`" section into the L1 instruction memories of the cores. This means that the definitions of the "guard" symbols are local to the processor. If the `.LDF` file is changed so that the "`cplb_code`" section is mapped into shared memory instead, then the `SHARED_MEMORY` directive must also define appropriate "guard" symbols, otherwise the link may resolve the reference by importing the default guard symbols from the run-time library.

Run-Time Library Functions

The three approaches discussed here are concerned primarily with arrangement of application code and data, but it is a rare application that does not make use of run-time library support in some manner. This raises complications for a dual-core system.

Reentrancy

The majority of run-time library routines make no use of private data, operating on parameters and stack data only. Such functions are fully usable within a dual-core system without the need for locking. Some Standard routines—such as `strtok()`—make use private data, and some routines update global data—the `errno` variable being the most common global variable so effected.

Multi-core applications must be built with the `-multicore` compiler switch, which means that the multicore variants of these functions will be used. They have the appropriate locking enabled, and allocate per-core private copies of such data to ensure that each core sees standardized behavior.

However, not all run-time library functions may be freely mapped. There are some restrictions on mapping. These are documented in "Library Function Re-Entrancy and Multi-Threaded Environments" on page 3-9.

Placement

In the majority of cases, run-time library code is placed into the `program` section, and data items are placed into a `data1`, `constdata` or `bsz` section depending on characteristics. This generic placement makes it difficult to partition library routines among the four major components of core A, core B, shared internal memory, and shared external memory.

Library code can be moved around on a coarse level by mapping specific objects to memory areas. A typical LDF mapping looks like:

```
INPUT_SECTIONS( $LIBRARIES(program) )
```

which expands the `$LIBRARIES` LDF macro into the list of library files, and then extracts "`program`" section objects from those libraries to resolve symbols. You can explicitly select particular objects from a library using the following method:

```
INPUT_SECTIONS( libc532.dlb [ printf.doj ] (program) )
```

This explicitly retrieves the "`program`" section items from the `printf.doj` object file within the `libc532.dlb` library, and maps it to the memory area corresponding to the `INPUT_SECTIONS` directive (note that spacing is significant in this syntax. See `INPUT_SECTIONS` in the `SECTIONS{}` directive section of the *VisualDSP++ 4.0 Linker and Utilities Manual*.)

This approach can be used to place library elements more precisely. The Expert Linker also provides a graphical environment for creating such mappings.

For restrictions on placing library functions in memory, see sections "Library Placement" on page 3-12 and "Section Placement" on page 3-13.

Restrictions On Dual-Core Applications

There are some restrictions for dual-core applications that do not apply to other applications.

Compiler Facilities

The following compiler facilities are not currently supported for the single/dual approach:

- Interprocedural Analysis (IPA) optimization (the `-ipa` compiler switch).
- Profile-Guided Optimization (PGO) (the `-pguide` compiler switch).
- Instrumented-code Profiling (`-p`, `-p1` and `-p2` compiler switches).

Cross-Core Memory References

It is not valid for code executing in one core to access the L1 memory of the other core, whether for code or data references. Attempts to do so will raise an exception. Therefore, when pointers to L1 memory are stored in shared memory and accessed by common code, care must be taken to ensure that such pointers are not dereferenced by the other core. This applies to both the per-core and single/dual approaches.

Debugger Breakpoints

Breakpoints placed in shared memory can cause problems when using the emulator. While the breakpoint is conceptually associated with a particular core, in reality the “emulator trap” instruction used affects both cores. This can lead to a core stopping unexpectedly, since the breakpoint is associated with the other core, or to a core running freely past a breakpoint because the “emulator trap” instruction has been temporarily removed while the other core was being stepped over the breakpoint.

Dual-Core Programming Examples

The following examples show the different code design approaches as applied to a simple client-server application on the ADSP-BF561 processor. The client passes a list of sentences to a server, one by one, and the server encodes them via a trivial `ROT13` algorithm. The client shows each string before encoding, after encoding, and once more after re-encoding (which, under `ROT13`, restores the original plain text). The contents of the buffer can be viewed in a memory window to investigate the client's progress.

A "`frame`" object is used to pass each sentence between the client and the server, and to return the encoded form.

Single-Core Application Example

The single-core approach is a single-threaded version, for simplicity. An alternative would be to use a VDK implementation that provided a thread for each of the client and the server.

Since there is just a single thread, no synchronization is necessary, so the main loop looks like this:

```
#include "defs.h
int main(void)
{
   int i;
   for (i = 0; frames[i]; i++) {
        int n = 0, m;
      char *s = frame.buffer;
      const char *t = frames[i];
      while (*s++ = *t++)
              n++;
      frame.len = n;
      report(PlainText, (char *)frame.buffer);
      rot13(frame.buffer);
      report(Encoded, (char *)frame.buffer);
```

```
      rot13(frame.buffer);
      report(Decoded, (char *)frame.buffer);
   }
   return 0;
}
```

This can be built using the default compiler `.LDF` file:

```
ccblkfn -o singlecore.doj -c -proc ADSP-BF561 -O singlecore.c
ccblkfn -o rot13.doj -c -proc ADSP-BF561 -O rot13.c
ccblkfn -o frame.doj -c -proc ADSP-BF561 -O - frame.c
ccblkfn -o corea.dxe -proc ADSP-BF561 -O singlecore.doj
           rot13.doj frame.doj report.doj
```

Because the "frame buffer" is only used by a single thread in a single core, it need not be placed into any particular data area. The client's output text is written into `output_buffer[]`:

```
typedef enum { PlainText, Encoded, Decoded } encoding;
char output_buffer[2000];

void report(encoding which, char *text)
{
    static int i;
    char *label;
    switch (which) {
    case PlainText: label = "Before: "; break;
    case Encoded:   label = "After:  "; break;
    default:        label = "Done:   "; break;
    }
    do {
        output_buffer[i++] = *label;
        i %= 2000;
    } while (*label++);
    do {
        output_buffer[i++] = *text;
        i %= 2000;
    } while (*text++);
}
```

One Application Per Core Example

The per-core approach requires some measure of synchronization, so locking routines are added to the build. It also requires another thread that runs in the second core. The two threads utilize a lock to serialize access to the buffer, and a protocol to indicate the buffer state:

```
#include <ccblkfn.h>
typedef enum { WaitingToBeProcessed, ProcessingDone, NoMoreWork }
BufState;
typedef struct {
        BufState state;
        char buffer[80];
} Frame;
extern volatile Frame frame; Frame;
```

- The buffer starts in state `ProcessingDone`—there is no work pending
- Core A copies data into the buffer and sets the state to `WaitingToBeProcessed`
- The buffer now belongs to core B which does the necessary encoding and resets the state to `ProcessingDone`
- The buffer now belongs to core A again, and core A is free to examine the results
- When Core A has passed all packets of data to Core B and received all the responses, Core A sets the state to `NoMoreWork`, to indicate to Core B that it can terminate.

Therefore, the main loop in core A (the “client” core) looks like this:

```
void wait_for_b(void)
{
   BufState state;
   do {
      claim_lock(&corelock, CoreA);
```

```
        state = frame.state;
        release_lock(&corelock);
        } while (state != ProcessingDone);
}

int main(void)
{
        int i;
        for (i = 0; frames[i]; i++) {
            int n = 0, m;
            char *s = (char *)frame.buffer;
            const char *t = frames[i];
            claim_lock(&corelock, CoreA);
            while (*s++ = *t++)
                    ;
            frame.state = WaitingToBeProcessed;
            report(PlainText, (char *)frame.buffer);
            release_lock(&corelock);

            wait_for_b();
            report(Encoded, (char *)frame.buffer);

            claim_lock(&corelock, CoreA);
            frame.state = WaitingToBeProcessed;
            release_lock(&corelock);

            wait_for_b();
            report(Decoded, (char *)frame.buffer);
        }
        claim_lock(&corelock, CoreA);
        frame.state = NoMoreWork;
        release_lock(&corelock);
        return 0;
}
```

The “server”, in core B, is waiting on the buffer being passed over, so that it can encode the contents and return the results:

```
int main(void)
{
```

```
    BufState state;
    do {
        claim_lock(&corelock, CoreB);
        state = frame.state;
        if (state == WaitingToBeProcessed) {
            rot13((char *)frame.buffer);
            frame.state = ProcessingDone }
      release_lock(&corelock);
    } while (state != NoMoreWork);
}
```

The basic locking functions are expanded in this example, so that each lock records which core owns it at the time. This adds no particular benefit to the example, but demonstrates the use of common routines:

```
typedef enum { Free = 0, CoreA = 'A', CoreB = 'B' } Core;
typedef struct {
      testset_t flag;
      Core owner;
} Lock;

void claim_lock(volatile Lock *lockptr, Core who)
{
      int state = cli();
      adi_acquire_lock(&lockptr->flag);
      lockptr->owner = who;
      sti(state);
}

void release_lock(volatile Lock *lockptr)
{
      int state = cli();
      lockptr->owner = Free;
      adi_release_lock(&lockptr->flag);
      sti(state);
}
```

Both cores use these common routines, although each core has a private copy in this case. Since the two routines must share the same memory location for the lock, a customized .LDF file is used. This is based on the default .LDF file, ADSP-BF561.ldf, with some additional mappings.

The shared data is explicitly mapped to sections. The objects are declared in core A.

```
section("sdram_shared") volatile Frame frame = { ProcessingDone };
section("l2_shared") volatile Lock corelock;
```

We then make extern references to these within core B:

```
extern volatile Frame frame;
extern volatile Lock corelock;
```

As these symbols have no definition in core B, we will arrange for them to be resolved to the corresponding definitions in core A, thus making core A and core B use the same address for the shared resources. This is done by adding appropriate RESOLVE directives to the .LDF file , within the PROCESSOR directive for core B:

```
PROCESSOR p1
{
    OUTPUT( $COMMAND_LINE_OUTPUT_FILE )
    /* Following address must match start of MEM_B_L1_PROGRAM */
    RESOLVE(start,0xFF600000)
#ifdef IDDE_ARGS
    RESOLVE(___argv_string,ARGV_START)
#endif
    /* Add the following two lines to the Processor
       directives for core B */

    RESOLVE(_frame,$COMMAND_LINE_OUTPUT_DIRECTORY/corea.dxe)
    RESOLVE(_corelock,$COMMAND_LINE_OUTPUT_DIRECTORY/corea.dxe)
```

The two applications (one per core) are built individually. Core A must be built before core B as core B depends on the core A executable at the link stage:

```
ccblkfn -o rot13.doj -c -O -proc ADSP-BF561 -multicore rot13.c
ccblkfn -o frame.doj -c -O -proc ADSP-BF561 -multicore frame.c
ccblkfn -o lockfns.doj -c -O -proc ADSP-BF561 -multicore lockfns.c
ccblkfn -o lockdata.doj -c -O -proc ADSP-BF561 -multicore lockdata.c
ccblkfn -o client.doj -c -O -proc ADSP-BF561 -multicore client.c
ccblkfn -o report.doj -c -O -proc ADSP-BF561 -multicore report.c

ccblkfn -o corea.dxe -proc ADSP-BF561 -flags-link
        -MD__ADI_MULTICORE \
        rot13.doj frame.doj lockfns.doj lockdata.doj report.doj \
        -flags-link -MDCOREA client.doj -T ./percore.ldf

ccblkfn -o server.doj -c -O -proc ADSP-BF561 -multicore server.c

ccblkfn -o coreb.dxe -proc ADSP-BF561 -flags-link
        -MD__ADI_MULTICORE \
        rot13.doj frame.doj lockfns.doj lockdata.doj\
        -flags-link -MDCOREB server.doj -T ./percore.ldf
```

The result is two `.DXE` files, `corea.dxe` and `coreb.dxe`, that can be loaded onto the two cores and executed in parallel.

Single Application/Dual Core Example

The sources for the single/dual approach are the same. The differences are in how they are linked into a single application.

Five projects are created—the overall project (Codec) and the four sub-projects (`corea`, `coreb`, `sml2` and `sml3`). All five projects are opened at once, with project dependencies set so that Codec depends on the other four.

A custom .LDF file is created for the Codec project, but *no customizing is done.* The single/dual approach requires the more advanced features that are available in the Expert Linker's templates, but the C template is sufficient "as-is" for this example.

The files are added to the projects in the following manner:

```
corea: client.c report.c
coreb: server.c
sml2: lockfns.c lockdata.c
sml3: frame.c rot13.c
```

This division is arbitrary and is used simply to demonstrate placement within the different shared memories. The lockdata.c and frame.c files contain the shared symbols, while lockfns.c and rot13.c contain the shared code.

The projects must be configured to perform multi-core builds:

- For each of the sub-projects, add "-multicore" to the **Additional options** field of the **Compile** tab, under **Project Options**.
- For the Codec project, add "__ADI_MULTICORE" to the **Preprocessor Macro Definitions** field of the **LDF Preprocessing** category, on the **Link** tab, under **Project Options**.

The entire application is built using a single build process:

```
----------------Configuration: sml3 - Debug----------------
  ccblkfn -c .\frame.c -g -proc ADSP-BF561 -o .\Debug\frame.doj
  ccblkfn -c .\rot13.c -g -proc ADSP-BF561 -o .\Debug\rot13.doj
  ccblkfn -proc ADSP-BF561 -build-lib -o .\Debug\sml3.dlb
          .\Debug\frame.doj .\Debug\rot13.dojc

----------------Configuration: sml2 - Debug----------------
  ccblkfn -c .\lockdata.c -g -proc ADSP-BF561 -o .\Debug\lockdata.doj
  ccblkfn -c .\lockfns.c -g -proc ADSP-BF561 -o .\Debug\lockfns.doj
```

```
  ccblkfn -proc ADSP-BF561 -build-lib -o
          .\Debug\sml2.dlb .\Debug\lockdata.doj
          .\Debug\lockfns.doj

----------------Configuration: coreb - Debug----------------
  ccblkfn -c .\server.c -g -proc ADSP-BF561 -o .\Debug\server.doj
  ccblkfn -proc ADSP-BF561 -build-lib -o .\Debug\coreb.dlb
          .\Debug\server.doj

----------------Configuration: corea - Debug----------------
  ccblkfn -c .\client.c -g -proc ADSP-BF561 -o .\Debug\client.doj
  ccblkfn -c .\report.c -g -proc ADSP-BF561 -o .\Debug\report.doj
  ccblkfn -proc ADSP-BF561 -build-lib -o .\Debug\corea.dlb
          .\Debug\client.doj. \Debug\report.doj

----------------Configuration: Codec  - Debug----------------
  ccblkfn -T .\dualcore.ldf -L .\Debug -flags-link -od,
          .\Debug -o .\Debug\Codec.dxe -proc ADSP-BF561
```

First, the sub-project libraries are built, and then the top-level project is used to build the whole application. The `.LDF` file specifies all the output files within it generating `p0.dxe`, `p1.dxe` and `sml2.sm`.

Shared Memory, Breakpoints and printf

When you load the built project into an EZ-KIT system using VisualDSP++, you may see warnings appear in the console:

```
Device 0:a: Breakpoints in shared memory are not recommended
because they may cause unexpected results.
```

VisualDSP++ warns about breakpoints in shared memory because breakpoints are conceptually private to a core, but physically apply to both cores. For example, if one core has been halted by a breakpoint, and is sin-

gle-stepping over it, VisualDSP++ must temporarily remove the breakpoint. If the other core is running freely, it may run past the breakpoint's location during that moment, without being stopped.

The message can be generated even if you have not placed any breakpoints in your project, because VisualDSP++ automatically plants breakpoints at particular symbols:

Breakpoint at symbol	Reason
`_main`	Supporting “Run to main”
`__lib_prog_term`	Halting if an application returns from `main()`, calls `exit()`, and so on.
`__primIO`	Used for host/target interaction, for File I/O. Required for `printf()` and other `stdio`-related functions.

When linking single/dual applications, care must be taken with File I/O operations. It is possible for both of the cores and the shared memory to resolve the File I/O library routines separately, leading to independent copies of the functions in the cores and in shared memory. This is not normally a problem, but can cause difficulties in this case because VisualDSP++ will only place a breakpoint into one instance of `__primIO`, and if the core is using the other copy, File I/O operations will not take place. If this happens when building the above example, both cores will run to completion successfully, but no output will appear in the console.

Synchronization Functions

VisualDSP++ 4.0 provides functionality for synchronization. There are two compiler intrinsics (build-in functions) and three locking routines.

The compiler intrinsics are:

```
#include <ccblkfn.h>
int testset(char *);
void untestset(char *);
```

The `testset()` intrinsic generates a native `TESTSET` instruction, which can perform atomic updates on a memory location. The intrinsic returns the result of the `CC` flag produced by the `TESTSET` instruction. Refer to the instruction set reference for details.

The `untestset()` intrinsic clears the memory location set by the `testset()` intrinsic. This intrinsic is recommended in place of a normal memory write because the `untestset()` intrinsic acts as a stronger barrier to code movement during optimization.

The three locking routines are:

```
#include <ccblkfn.h>
void adi_acquire_lock(testset_t *);
int adi_try_lock(testset_t *);
void adi_release_lock(testset_t *);
```

The `adi_acquire_lock()` routine repeatedly attempts to claim the lock by issuing `testset()` until successful, whereupon it returns to the caller. In contrast, the `adi_try_lock()` routine makes a single attempt—if it successfully claims the lock, it returns nonzero, otherwise it returns zero.

The `adi_release_lock()` routine releases the lock obtained by either `adi_acquire_lock()` or `adi_try_lock()`. It assumes that the lock was already claimed and makes no attempt to verify that its caller is in fact the current owner of the lock.

None of these intrinsics or functions disable interrupts—that is left to the caller's discretion.

I INDEX

Symbols

µ-law compression function, 4-128
µ-law expansion function, 4-129

Numerics

16-bit fractional ETSI routines, 1-123
2-d convolution (conv2d3x3) function, 4-81
2-d convolution (conv2d) function, 4-79
32-bit fractional ETSI routines, 1-119
32-bit saturation, 1-53
40-bit saturation, 1-53
64-bit counter, 4-42

A

-A (assert) compiler switch, 1-25
Abend (See abort function)
abort (abnormal program end) function, 3-54
abs (absolute value) function, 3-55
absolute value (see abs, fabs, labs functions)
accumulator, 1-53
accumulators, 1-208
a_compress (A-law compression) function, 4-44
acos (arc cosine) function, 3-56
acosd function, 3-56
acosf function, 3-56
acos_fr16 function, 3-56
add_devtab_entry function, 3-40
add_devtab_entry() function, 3-40
address
  event vector table, 1-235
adi_acquire_lock function, 3-57
adi_acquire_lock() routine, A-31
adi_core_id function, 3-59
_ADI_FAST_ETSI macro, 1-116
adi_free_mc_slot function, 3-60
adi_get_mc_value function, 3-60
__ADI_LIBEH__ macro, 1-68
_ADI_LIBIO macro, 1-43
__ADI_MULTICORE macro, 1-40, 3-11
adi_obtain_mc_slot function, 3-60
adi_release_lock function, 3-57
adi_release_lock() routine, A-32
adi_set_mc_value function, 3-60
__ADI_THREADS macro, 1-58
adi_try_lock function, 3-57
adi_try_lock() routine, A-31
ADSP-BF561 processor
  architecture overview, A-3, A-6
  internal memory, A-9
  L2 internal memory, A-6
  locking routines, A-31
  one-application-per-core, A-8
  run-time library support, A-18
  sharing code in L2 memory, A-16
  single application/dual core approach, A-13
  single-core application, A-5, A-6
  startup code, A-12
  synchronization functions, A-31
ADSP-BF561 run-time library routines
  code placement, A-18
  reentrancy, A-18
__ADSPBLACKFIN__ macro, 1-226

__ADSPBLACKFIN__ preprocessor
macro, 1-51
__ADSPLPBLACKFIN__ macro, 1-226
a_expand (A-law expansion) function, 4-45
aggregate
initializer, 1-100
return pointer, 1-252
A-law
compression, 4-44
expansion, 4-45
algebraic functions (See math functions)
alias
avoiding, 2-16
alignment
data, 1-154
alignment inquiry keyword, 1-194
__alignof__ (type-name) construct, 1-194
ALLDATA qualifier, 1-162
alloca() function, 1-137
allocate memory (See calloc, free, malloc, realloc functions)
all workaround, 1-62
alog10 functions, 4-48
alphanumeric character test (See isalnum function)
alternate keywords, 1-42
alternative tokens
disabling, 1-41
-alttok (alternative tokens) C++ mode compiler switch, 1-26
-anach (enable C++ anachronisms) C++ mode compiler switch, 1-67
anachronisms
default C++ mode, 1-67
disabled C++ mode, 1-69
__ANALOG_EXTENSIONS__ macro, 1-226
annotations
assembly code, 2-54
assembly source code position, 2-63
loop identification, 2-58
vectorization, 2-71
anomaly 05-00-0048, 1-63
anomaly 05-00-0054, 1-64
anomaly 05-00-0071, 1-63, 1-203
anomaly 05-00-0087, 1-63
anomaly 05-00-0104, 1-62
anomaly 05-00-0109, 1-236
anomaly 05-00-0127, 1-64
anomaly 05-00-0157, 1-65
anomaly 05-00-0164, 1-64
anomaly 05-00-0165, 1-64
anomaly 05-00-0195, 1-65
anomaly 05-00-0198, 1-65
anomaly workarounds, 1-62
ANSI C signal handler, 1-204
ANSI/ISO standard C++, 1-24
ANSI standard warnings, 1-49
anti-log
functions, 4-46
application binary interface, 1-72
arc cosine, 3-56
arc sine, 3-65
arc tangent, 3-66
arc tangent of quotient, 3-67
argc parameter, 1-243
argc support, 1-197
arg (get phase of a complex number) function, 4-50
argument
and return transfer, 1-257
passing, 1-257
argv/argc
arguments, 1-197
__argv global array, 1-243
argv parameter, 1-243

___argv_string variable
  defining, 1-197
__argv_string variable, 1-197
argv support, 1-197
arithmetic functions, 4-5
array
  sorting, 3-167
  zero length, 1-191
array search, binary (See bsearch function)
ASCII string (*see* atof, atoi, atol, atold functions)
ASCII string (See atof, atoi, atol functions)
asctime (convert broken-down time into a string) function, 3-63
asctime function, 3-26, 3-98
asin (arc sine) function, 3-65
asind function, 3-65
asinf function, 3-65
asin_fr16 function, 3-65
asm
  keyword, 1-194
  statement, 1-192, 2-20
asm()
  compiler keyword, 1-79
    (See also Assembly language support keyword (asm))
  constructs
    registers for, 1-87
    reordering, 1-94
  operand constraints, 1-88, 1-91
  template in C programs, 1-84
asm() construct, 1-82
  operands, 1-87
asm() constructs
  flow control, 1-95
  syntax, 1-84
  syntax rules, 1-85
assembler
  Blackfin processors, 1-3
assembly
  code annotations, 2-54
assembly language support keyword (asm), 1-82
asynchronous data change, 1-223
atan2 (arc tangent of quotient) function, 3-67
atan2d function, 3-67
atan2f function, 3-67
atan2_fr16 function, 3-67
atan (arc tangent) function, 3-66
atand function, 3-66
atanf function, 3-66
atan_fr16 function, 3-66
atexit function, 3-9
atexit (select exit function) function, 3-68
atof (convert string to double) function, 3-69
atoi (convert string to integer) function, 3-72
atol (convert string to long integer) function, 3-73
atold (convert string to long double) function, 3-74
atoll (convert string to long long integer) function, 3-77
__attribute__ keyword, 1-195
attributes
  functions, variables and types, 1-195
autocoh (autocoherence) function, 4-51
autocoherence, 4-51
autocorr (autocorrelation) function, 4-53
autocorrelation, 4-53
automatic
  inlining, 1-46, 1-73, 2-19
  loop control variables, 2-31
avoid-dag1 workaround, 1-62
avoid-dag-load-reuse workaround, 1-63

B

bank qualifier, 1-96, 2-23, 2-50
bank(string)
  compiler keyword, 1-79
base 10
  anti-log functions, 4-48
  logarithms, 3-154
basic complex arithmetic functions, 4-5
basiccrt.s file, 1-233
basic cycle counting, 4-33
big-endian, 1-136
binary array search (See bsearch function)
bit field, 2-11
bitfields
  signed, 1-55
  unsigned, 1-59
BITS_PER_WORD constant, 3-45
Blackfin-specific functionality, 1-196
  argv/argc arguments, 1-197
  caching of external memory, 1-208
  heap size, 1-201
  interrupts, 1-201
  profiling for single-threaded systems, 1-197
  profiling routine, 1-197
    for single-threaded systems, 1-197
Blackman-Harris window, 4-104
blank space character, 3-142
Boolean type support keywords (bool, true, false), 1-98
breakpoint
  debugger, A-20
broken-down time, 3-24, 3-151, 3-193
bsearch (binary search in sorted array) function, 3-78
-bss compiler switch, 1-27
BSZ qualifier, 1-162
buf field, 3-47
build-in functions
  synchronization, A-31
-build-lib (build library) compiler switch, 1-27
__builtin_aligned function, 2-7, 2-15, 2-49
__builtin_circindex function, 2-38, 2-39
__builtin_circptr function, 2-38, 2-39
built-in functions, 1-103
  circular buffers, 1-134
  _clip, 1-105
  compiler, 3-5
  complex fract, 1-130
  endian swapping, 1-137
  ETSI, 1-105, 1-115
  exceptions, 1-138
  fract16, 1-104, 1-106
  fract2x16, 1-104, 1-110
  fract32, 1-104, 1-108
  fractional arithmetic, 1-104
  fracts in C, 1-129
  fracts in C++, 1-128
  IMASK, 1-138
  in code optimization, 2-35
  interrupts, 1-138
  math, standard, 3-5
  misaligned data, 1-147
  synchronization, 1-138
  system, 1-137, 2-37
  testset(), A-31
  untestset(), A-31
  video operations, 1-140
  Viterbi functions, 1-132
__builtin prefix, 1-104
byteswap2, 1-136
byteswap4, 1-136

C

C
  fractional literal values, 1-129
  variable-length arrays, 1-102

C++
Abridged Library, 3-27
complex class, 1-131
constructor invocation, 1-242
fractional classes, 1-128
gcc compatibility features not supported, 1-187
streambuf class, 3-34
template inclusion control pragma, 1-176
virtual lookup tables, 1-54
-c89 (ISO/IEC 9899
1990 standard) compiler switch, 1-24
cabs (complex absolute value) function, 4-55
cache
asynchronous change systems, 1-223
changing configuration, 1-214
configuration definition, 1-208
configurations, 1-212
default configuration, 1-213
enabling, 1-212
enabling on ADSP-BF535 processor, 1-212
flushing, 1-220, 3-108
invalidating, 1-214, 3-80
modes, 1-218
protection lookaside buffers (CPLB), 1-28
protection lookaside buffers (CPLBs), 1-208, 1-209, A-7
write-back, 1-64
cache_invalidate function, 3-80
cache_invalidate routine, 1-214, 3-80, 3-91
caching
external memory, 1-208
cadd (complex addition) function, 4-56
calendar time, 3-24, 3-224
calling
assembly language subroutine, 1-269
library functions, 3-3
CALL instruction, 1-172
calloc (allocate and initialize memory) function, 3-83
call preserved registers, 1-252
C and C++ library files, 3-5
Carry
global variable, 1-117
Carry flag
for ETSI functions, 1-117
cartesian (Cartesian to polar) function, 4-57
case label, 1-193
case-sensitive switches, 1-5
C/assembly interfacing (See mixed C/assembly programming)
C/C++
code optimization, 2-2
language extensions, 1-78
library function calling, 3-3
run-time header (CRT), 1-209, 1-232
run-time model, 1-230
ccblkfn (Blackfin C/C++ compiler), 1-1, 1-2
C/C++ compiler
guide, 1-1, 1-2
overview, 1-1, 1-2
C/C++ compiler mode switches
-c89, 1-24
-c++ (C++ mode), 1-24
C/C++ language extensions
asm keyword, 1-82
bool keyword, 1-80
false keyword, 1-80
inline keyword, 1-81
restrict, 1-80
section() keyword, 1-80
true keyword, 1-80

-C (comments) compiler switch, 1-27
-c (compile only) compiler switch, 1-27
C/C++ run-time
  environment
    (See also mixed C/C++/assembly programming)
  environment, defined, 1-230
C/C++ run-time libraries, 3-3
  linking, 3-5
  start-up files, 3-8
  thread-safe, 3-34
  variants, 3-5, 3-7
C/C++ run-time library, 3-5
  files, 3-5
C data types, 1-261
cdiv (complex division) function, 4-59
ceil (ceiling) functions, 3-84
cexp (complex exponential) function, 4-60
cfft2d (n x n point 2-d complex input FFT) function, 4-68
cfftf (fast N-point radix-4 complex input FFT) function, 4-63
cfft (n point radix-2 complex FFT) function, 4-61
cfftrad4 (n point radix-4 complex FFT) function, 4-66
cfir (complex FIR filter) function, 4-70
char
  data type, 1-261
 character, 1-40, 1-44
character string search (See strchr function)
char storage format, 1-261
circular
  buffer registers, 1-208
  indexing, 1-134
circular buffers, 1-251, 2-38
  automatic generation, 1-134
  explicit circular buffer generation, 1-135
  generating, 1-134
  increment of index., 1-135
  increments of pointer, 1-135
  set to zero, 1-236
  switch setting, 1-32
C language extensions
  C++ style comments, 1-80
  indexed initializers, 1-80
  non-constant initializers, 1-80
  preprocessor-generated warnings, 1-80
  variable length arrays, 1-80
C library facilities, 3-30
cli() function, 1-138
CLI instruction, 1-203
_clip built-in functions, 1-105
clip (clip) function, 4-72
clobber, 1-85
clobbered
  registers, 1-85, 1-169, 1-171
  register sets, 1-171
clobber string specifiers, 1-90
clock function, 4-38, 4-41
clock (processor time) function, 3-85
clock speed, 1-238
CLOCKS_PER_SEC macro, 3-25, 4-38, 4-40
clock_t data type, 3-24, 3-85
close function, 3-37
cmlt (complex multiply) function, 4-73
C++ mode compiler switches
  -anach (enable C++ anachronisms), 1-67
  -eh (enable exception handling), 1-68
  -no-anach (disable C++ anachronisms), 1-69
  -no-demangle (disable demangler), 1-69
  -no-eh (disable exception handling), 1-69
  -no-implicit-inclusion, 1-69
  -no-rtti (disable run-time type identification), 1-69
  -rtti (enable run-time type identification), 1-70

CODE_FAULT_ADDR, 1-206
CODE_FAULT_STATUS, 1-206
code optimization
  enabling, 1-45
  for size, 1-47
CODE qualifier, 1-162
coeff_iirdf1_fr16 function, 4-74
coeff_iirdf1 function, 4-74
command-line
  interface, 1-4 to 1-75
  syntax, 1-5
comparing
  characters in strings, 3-199
  null-terminated strings, 3-188
compiler
  built-in functions, 3-5
  C/C++ language extensions, 1-78
  code generator workarounds, 1-62
  code optimization, 1-71, 2-2
  command-line switch summaries, 1-13
  command-line syntax, 1-5
  input/output files, 1-7
  optimizer, 2-4
  overview, 1-2
  performance functions, 1-139
  profiling facilities, 1-197
  running, 1-5
  specifying functional options, 1-9
compiler common switches
  -A (assert), 1-25
  -bss, 1-27
  -build-lib (build library), 1-27
  -C (comments), 1-27
  -c (compile only), 1-27
  -const-read-write, 1-27
  -cplbs (CPLBs are active), 1-28
  -D (define macro), 1-28
  -debug-types, 1-28
  -decls, 1-29
  -default-linkage, 1-29
  -double-size-{32 | 64}, 1-29
  -double-size-any, 1-29
  -dry (a verbose dry-run), 1-30
  -dryrun (a terse dry-run), 1-30
  -ED (run after preprocessing to file), 1-30
  -EE (run after preprocessing), 1-30
  -E (stop after preprocessing), 1-30
  -extra-keywords (enable short-form keywords), 1-31
  -extra-loop-loads, 1-31
  -fast-fp (fast floating point), 1-31
  -@ filename, 1-25
  -flags (command-line input), 1-31
  -force-circbuf, 1-32
  -force-link, 1-32
  -fp-associative (floating-point associative operation), 1-32
  -full-io, 1-33
  -full-version (display version), 1-33
  -g (generate debug information), 1-33
  -guard-vol-loads, 1-34
  -help (command-line help), 1-34
  -HH (list headers and compile), 1-34
  -H (list headers), 1-34
  -I directory (include search directory), 1-34
  -I- directory (start include directory list), 1-35
  -ieee-fp (slow floating point), 1-36
  -i (less includes), 1-36
  -include (include file), 1-36
  -ipa (interprocedural analysis), 1-36
  -jcs21, 1-37
  -jcs21+, 1-37
  -jump-constdata|data|code, 1-37
  -L (library search directory), 1-37
  -l (link library), 1-38
  -map (generate a memory map), 1-40

-MD (generate make rules and compile), 1-39
-mem (invoke memory initializer), 1-40
-M (generate make rules only), 1-38
-MM (generate make rules and compile), 1-39
-Mo (processor output file), 1-39
-MQ (output without compilation), 1-39
-Mt (output make rule for named file), 1-39
-multicore, 1-40
-multiline, 1-40
-no-alttok (disable alternative tokens), 1-41
-no-annotate (disable alternative tokens), 1-41
-no-bss, 1-41
-no-builtin (no built-in functions), 1-41
-no-circbuf (no circular buffer), 1-42
-no-defs (disable defaults), 1-42
-no-extra-keywords, 1-42
-no-force-link, 1-42
-no-fp-associative, 1-43
-no-full-io, 1-43
-no-int-to-fact (disable integer to fractional conversion), 1-43
-no-int-to-fract, 1-43
-no-jcs21+, 1-44
-no-jcs2l, 1-44
-no-mem (not invoking memory initializer), 1-44
-no-multiline, 1-44
-no-saturation (no faster operations), 1-44
-no-std-ass (disable standard assertions), 1-44
-no-std-def (disable standard macro definitions), 1-45
-no-std-inc (disable standard include search), 1-45
-no-std-lib (disable standard library search), 1-45
-no-threads (disable thread-safe build), 1-45
-Oa (automatic function inlining), 1-46
-O (enable optimizations), 1-45
-Ofp (frame pointer optimizations), 1-46
-Og (optimize while preserving debugging information), 1-46
-o (output file), 1-47
-Os (enable code size optimizations), 1-47
-Ov (optimize for speed vs. size), 1-47
-path-install (installation location), 1-48
-path-output (non-temporary files location), 1-48
-path-temp (temporary files location), 1-48
-path (tool location), 1-48
-pchdir directory (locate precompiled header repository), 1-49
-pch (recompiled header), 1-49
-pedantic (ANSI standard warnings), 1-49
-pedantic-errors (ANSI standard errors), 1-49
-p (generate profiling implementation), 1-47
-pguide (profile-guided optimization), 1-50
-P omit line numbers), 1-47
-pplist (preprocessor listing), 1-50
-PP (omit line numbers and run), 1-47
-progress-rep-func, 1-51
-progress-rep-gen-opt, 1-51
-progress-rep-mc-opt, 1-52
-R (add source directory), 1-52
-R- (disable source path), 1-52

-reserve (reserve register), 1-53
-sat32 (32 bit saturation), 1-53
-sat40 40 bit saturation), 1-53
-save-temps (save intermediate files), 1-54
-sdram, 1-54
-section (data placement), 1-54
-show (display command line), 1-55
-signed-bitfield (make plain bitfields signed), 1-55
-signed-char (make char signed), 1-55
-si-revision version (silicon revision), 1-55
sourcefile, 1-25
-S (stop after compilation), 1-53
-s (strip debug information), 1-53
-syntax-only (only check syntax), 1-57
-sysdef (system definitions), 1-57
-threads (enable thread-safe build), 1-58
-time (tell time), 1-59
-T (linker description file), 1-58
-unsigned-bitfield (make plain bitfields unsigned), 1-59
-unsigned-char (make char unsigned), 1-60
-U (undefine macro), 1-59
-verbose (display command line), 1-60
-version (display version), 1-60
-v (version and verbose), 1-60
-warn-protos (warn if incomplete prototype), 1-62
-w (disable all warnings), 1-61
-Werror-limit (maximum compiler errors), 1-61
-workaround (workaround id), 1-62
-W (override error message), 1-60
-Wremarks (enable diagnostic warnings), 1-61
-write-files (enable driver I/O redirection), 1-65
-write-opts (user options), 1-66
-Wterse (enable terse warnings), 1-61
-xref (cross-reference list), 1-66
complex
absolute value, 4-55
addition, 4-56
conjugate, 4-76
division, 4-59
exponential, 4-60
fract built-ins, 1-130
functions, 4-5
multiply, 4-73
number, 4-50
subtraction, 4-89
complex_fract16 type functions, 1-130
complex_fract32, 1-130
complex header file, 3-27
complex.h header file, 4-5
compression/expansion, 4-12
conditional code
in loops, 2-29
conditional expressions
with missing operands, 1-190
conj (complex conjugate) function, 4-76
const
keyword, 2-33
pointers, 1-27
qualifier, 2-33
constants
initializing statically, 2-13
CONSTDATA qualifier, 1-162
constraint, 1-84
-const-read-write
compiler switch, 1-28
compile-time flag, 2-33

constructs
flow control, 1-95
input and output operands, 1-94
operand description, 1-87
reordering and optimization, 1-94
template for assembly, 1-84
with multiple instructions, 1-93
continuation characters, 1-40, 1-44
control character test (See iscntrl function)
controlling
optimization, 1-71
conv2d (2-d convolution) function, 4-79
conv2d3x3 (2-d convolution) function, 4-81
conversion specifiers, 3-22
convert
characters (See tolower, toupper functions)
implicit type, 3-20
strings (See atof, atoi, atol,strtok, strtol, strtoul, functions)
convolution, 4-8
convolve (convolution) function, 4-77
copying
characters from one string to another, 3-200
from one string to another, 3-190
copysign (copysign) function, 4-82
core
identifying current, 3-57, 3-59
cos (cosine) function, 3-86
cosd function, 3-86
cosf function, 3-86
cos_fr16 function, 3-86
coshd function, 3-88
coshf function, 3-88
cosh (hyperbolic cosine) functions, 3-88
cosine, 3-86
cotangent, 4-83
cot (cotangent) function, 4-83
counting one bits in word, 4-84
countones (count one bits in word) function, 4-84
count_ticks() function, 1-184
CPLB
initialization, 1-240, 3-91
cplb_code memory section, 1-211
___cplb_ctrl, 3-80
___cplb_ctrl control variable, 1-209, 1-211, 1-212, 1-214, 1-216, 1-222, A-12
___cplb_ctrl variable, 1-236, 3-91
cplb_data memory section, 1-211
CPLB_ENABLE_CPLBS macro, 1-209
CPLB_ENABLE_DCACHE2 macro, 1-210
CPLB_ENABLE_DCACHE macro, 1-210
CPLB_ENABLE_ICACHE macro, 1-209
CPLB eviction, 3-94
CPLB exception handler, 3-89
___cplb_hdr default handler, 1-236
cplb_hdr function, 3-89
cplb_hdr routine, 3-89
cplb.h header file, 1-209
cplb_init function, 3-91
_cplb_init routine, 1-211
cplb_init routine, 3-91
cplb_mgr function, 3-94
_cplb_mgr routine, 1-211, 1-218
definition, 1-221
error codes returned, 1-222
flushing writes, 1-219
return codes, 1-221
cplb_mgr routine, 3-94
CPLB_NO_ADDR_MATCH return code, 1-222
CPLB_NO_UNLOCKED return code, 1-222
CPLB_PROT_VIOL return code, 1-222

CPLB_RELOADED return code, 1-222
CPLBs, 1-208
  enabling, 1-240
  installation, 1-211
-cplbs (CPLBs are active) compiler switch, 1-28
CPLB_SET_DCBS macro, 1-210
CPLUSCRT macro, 1-242
__cplusplus macro, 1-226
crosscoh (cross-coherence) function, 4-85
crosscorr (cross-correlation) function, 4-87
CRT
  header, 1-232
crt*.doj startup files, 3-7
C run-time
  library reference, 3-53 to 3-231
csub (complex subtraction) function, 4-89
csync() function, 1-138
CSYNC instruction, 1-63
  avoid silicon anomalies, 1-63
csync workaround, 1-63
ctime (convert calendar time into a string) function, 3-98
ctime function, 3-63
___ctorloop function, 1-242
___ctorloop library function, 1-242
C-type functions
  iscntrl, 3-132
  isgraph, 3-134
  islower, 3-137
  isprint, 3-140
  ispunct, 3-141
  isspace, 3-142
  isupper, 3-143
  isxdigit, 3-144
  tolower, 3-225
  toupper, 3-226
ctype.h header file, 3-49
cycle
  counter, 1-200
  counter, enabling, 1-238
  counts, 1-200
cycle count, 4-7
  computing, 1-200
  interrupt dispatcher, 1-208
cycle counter
  reading, 2-37
cycle_count.h header file, 4-7, 4-33, 4-34
cycle count register, 4-33, 4-41, 4-42
CYCLES2 register, 4-42
cycles.h header file, 3-25, 4-8, 4-35, 4-36
CYCLES_INIT(S) macro, 4-35
CYCLES_PRINT(S) macro, 4-36
CYCLES register, 4-42
CYCLES_RESET(S) macro, 4-36
CYCLES_START(S) macro, 4-35
CYCLES_STOP(S) macro, 4-35
cycle_t data type, 4-34

D

DAG port
  selecting, 1-238
DAG registers, 1-251
data
  alignment pragmas, 1-149, 1-150
  fetching, with 32-bit loads., 2-14
  field, 3-36
  packing, 3-45
  storage formats, 1-261
  word alignment, 2-14
data alignment
  misaligned accesses, 1-147, 1-154
data cache
  disabling, 3-100
  enabling, 3-102
  flushing, 3-108
DATA_FAULT_ADDR, 1-206
DATA_FAULT_STATUS, 1-206

DATA qualifier, 1-162
data type
  formats, 1-261
  scalar, 2-8
  sizes, 1-261
date information, 3-24
__DATE__ macro, 1-226
Daylight Saving flag, 3-24
dcache_invalidate_both routine, 1-215, 3-80
dcache_invalidate routine, 1-215, 3-80
-DCLOCKS_PER_SEC= compile-time switch, 4-40
-D (define macro) compiler switch, 1-28, 1-59
-DDO_CYCLE_COUNTS compile-time switch, 4-35, 4-41
-DDO_CYCLE_COUNTS switch, 4-34
deallocate memory (See free function)
debugging
  source-level, 1-46
  sourcr-level, 1-33
debug information
  removing, 1-53
-debug-types compiler switch, 1-28
declarations
  mixed with code, 1-193
-decls compiler switch, 1-29
dedicated registers, 1-251
default
  cache configuration tables, 1-213
  device, 3-43
  environment, 1-196
  handler, 1-204
  heap, 1-245
  LDF file, 1-201
  sections, 1-162
  startup code, 1-196, 1-211
  target processor, 1-51
default enable mask, 1-237
default I/O run-time library, 3-21
-default-linkage (default linkage type) compiler switch, 1-29
delete operator
  with multiple heaps, 1-250
demangler, 1-69
detecting control character, 3-132
detecting uppercase characters, 3-143
detect punctuation character (ispunct) function, 3-141
DevEntry structure, 3-35
device
  default, 3-43
  driver, 3-35
  drivers, 3-35
  identifiers, 3-35
  initialization, 1-240
  pre-registering, 3-41
device.h header file, 3-16, 3-35
DeviceID field, 3-36
device_int.h header file, 3-16
devtab.c library source file, 3-41
diagnostics
  control pragma, 1-178
difftime (difference between two calendar times) function, 3-99
DIRTY flag, 1-219
disable_data_cache function, 3-100
disable_data_cache routine, 3-100
div (division) function, 3-101
division
  complex, 4-59
division (See div, ldiv functions)
DMA, 1-253
double
  data type, 1-261, 1-262, 1-264
  data type formats, 1-29
  representation, 3-205
  storage format, 1-261
DOUBLE32 qualifier, 1-163

DOUBLE64 qualifier, 1-163
DOUBLEANY qualifier, 1-163
__DOUBLES_ARE_FLOATS__ macro, 1-226
-double-size-{32 | 64} compiler switch, 1-29
-double-size-32 compiler switch, 1-29, 1-261
-double-size-32 option, 1-263
-double-size-64 compiler switch, 1-29, 1-261, 1-262
-double-size-64 option, 1-263
-double-size-any compiler switch, 1-29, 1-261, 1-264
-dry-run (verbose dry-run) compiler switch, 1-30
-dry (terse -dry-run) compiler switch, 1-30
DSP
  filters, 4-8
  header files, 4-4
  run-time library format, 4-43
DSP run-time library
  calling function in, 4-3
  linking functions, 4-2
  source code, 4-3
  variants, 4-2
dual-core applications
  restrictions for, A-20
dual-core processor, A-1

E

easmblkfn assembler, 1-3
__ECC__ macro, 1-226
__EDG__ macro, 1-226
__EDG_VERSION__ macro, 1-226
-ED (run after preprocessing to file) compiler switch, 1-30
-EE (run after preprocessing) compiler switch, 1-30
-eh (enable exception handling) C++ mode compiler switch, 1-68
elfar archive library, 1-3
embedded C++ header files
  complex.h, 3-27
  exception.h, 3-28
  fract.h, 3-28
  fstream.h, 3-28
  iomanip.h, 3-28
  iosfwd.h, 3-28
  ios.h, 3-28
  iostream.h, 3-29
  istream.h, 3-29
  new.h, 3-29
  ostream.h, 3-29
  shortfract.h, 3-29
  sstream.h, 3-29
  stdexcept.h, 3-29
  streambuf.h, 3-30
  string.h, 3-30
  strstream.h, 3-30
Embedded C++ Library, 3-27
embedded standard template library, 3-31
emulated arithmetic
  avoiding, 2-12
enable_data_cache function, 3-102
enable_data_cache routine, 3-102
enabling
  optimization and debugging, 1-46
endian-swapping intrinsics, 1-136
End (See atexit, exit functions)
environment variables, 1-70
  ADI_DSP, 1-71
  CCBLKFN_OPTIONS, 1-71
  PATH, 1-70
  TEMP, 1-70
  TMP, 1-70
environment variablesCCBLKFN_IGNORE_ENV, 1-71

errno global variable, 3-9, 3-26
errno.h header file, 3-16
error messages
  control pragma, 1-178
escape character, 1-194
-E (stop after preprocessing) compiler switch, 1-30
ETSI
  built-in functions, 1-115, 1-118, 1-267
  routines for fracts, 1-128
  run-time support library, 3-6
  support routines, 1-116
ETSI library
  carry flag, 1-118
  overflow flag, 1-118
ETSI routines
  16-bit fractional, 1-123
  32-bit fractional, 1-119
  RND_MOD flag, 1-118
ETSI_SOURCE macro, 1-115, 1-117, 1-128
ETSI support routines
  in libetsi*.dlb library, 1-116
event
  context, gaining access to, 1-206
  handlers, 3-177
  handlers registering directly, 3-174
  vector table, 1-156, 1-204, 3-177
event details
  exceptions, 1-206
  fetching, 1-206
event handler, 1-202
Event Vector Table, 1-236
EXCAUSE values, 1-206
exception handler, 1-221
  calling _cplb_mgr, 1-211
  _cplb_hdr, 1-211
  CPLBs, 3-89
  disabling, 1-69
  enabling, 1-68
exception handling, 3-94
exception header file, 3-28
exception.h file, 1-204, 1-205
exceptions, 1-202
__EXCEPTIONS macro, 1-68, 1-226
executable
  running, 1-198
EX_EXCEPTION_HANDLER compiler macro, 1-202
EX_INT_DEFAULT value, 1-204
EX_INTERRUPT_HANDLER compiler macro, 1-202
EX_INT_IGNORE value, 1-204
_exit
  calling, 1-243
exit (normal program termination) function, 3-104
EX_NMI_HANDLER compiler macro, 1-202
expected_false built-in function, 1-139, 2-24
expected_true built-in function, 1-139, 2-24
exp (exponential) functions, 3-105
exponential (see exp, ldexp functions)
exponentiation, 4-46, 4-48
EX_REENTRANT_HANDLER compiler macro, 1-202
extension keywords, 1-79
external memory, 1-54
-extra-keywords (not quite -analog) compiler switch, 1-31
-extra-loop-loads compiler switch, 1-31
EZ-KIT Lite system, 3-16, 3-23, 3-35, 3-43
  ADSP-BF561 processor, A-3
  with 27 MHz on-board clock, 1-239

F

fabs (absolute value) functions, 3-106

false (See Boolean type support keywords (bool, true, false))
far jump return (See longjmp, setjmp functions)
Fast Fourier Transforms, 4-8, 4-9
-fast-fp compiler switch, 1-261
-fast-fp (fast floating point) compiler switch, 1-31, 1-267
fetching event details, 1-206
FFT function versions, 4-8
file
  annotation position, 2-63
  extension, 1-5
  extensions, 1-7, 1-25
  searching, 1-7
fileID field, 3-48
file I/O
  extending to new devices, 3-35
  support, 3-34
__FILE__ macro, 1-226
filename
  reading from, 1-25
-@ filename (command file) compiler switch, 1-25
File-to-Device stream, 1-73
filter.h header file, 4-8, 4-95
filter library, 4-8
filters
  signal processing, 4-8
finite impulse response (FIR) filter, 4-90
FIOCRT macro, 1-243
fir_decima (fir decimation filter) function, 4-92
fir (finite impulser response filter) function, 4-90
fir_interp (FIR interpolation filter) function, 4-94
fir_interp_fr16 function, 4-94
-flags (command line input) compiler switch, 1-31
flags field, 3-46
float
  data type, 1-261, 1-262
  storage format, 1-261
float.h header file, 3-17
floating-point
  data size, 1-262
  emulation library, 1-15, 1-31
  numbers, 1-261
floating-point constants
  hexadecimal, 1-191
floor (integral value) functions, 3-107
flow control operations, 1-95
FLT_MAX macro, 3-17
FLT_MIN macro, 3-17
flush_data_buffer routine, 3-108
flush_data_cache function, 3-108
flush_data_cache routine, 3-108
flush_data_cache routine., 1-220
flushing
  data cache, 3-108
fmod (floating-point modulus) functions, 3-110
fopen() function, 3-43
-force-circbuf (circular buffer) compiler switch, 1-32, 2-38
-force-link (force stack frame creation) compiler switch, 1-32
-fp-associative (floating-point associative) compiler switch, 1-32
fract
  data type, 1-261
fract16 built-in functions, 1-106
fract16 data type, 1-104, 1-267
fract2x16 built-in functions, 1-110
fract2x16 data type, 1-104, 1-267
fract32 built-in functions, 1-108
fract32 data type, 1-104, 1-267
fract class, 1-128

fract functions
ETSI, 1-115
fract header file, 3-28
fract.h header file, 1-105, 1-115
fractional
built-in functions, 1-104
built-in values, 1-104
complex_fract16 values, 1-130
C type values, 1-104
data,, 2-35
fract class values, 1-128
literal values in C, 1-129
shortfract class values, 1-128
fractional number built-in functions, 1-267
fractional numbers, 1-267
fract_math.h header file, 1-115
fract_math.h heafer file, 1-116
frame pointer, 1-46, 1-207, 1-237, 1-253, 1-254, 1-255
dedicated register, 1-251
fread() function, 3-22
free (deallocate memory) function, 3-111
free-list, 1-250
frexp (separate fraction and exponent) function, 3-112
fstream header file, 3-28
fstreams.h header file, 3-33
-full-io compiler switch, 1-33
-full-version (display version) compiler switch, 1-33
function
synchronization, 3-57
function pragmas, 2-42
functions
arguments/return value transfer, 1-257
arithmetic, 4-5
calling in loop, 2-30
complex, 4-5
entry (prologue), 1-253
exit (epilogue), 1-253
inlining, 2-19
math, 4-12
matrix, 4-16
statistical, 4-22
transformational, 4-8
vector, 4-24
function side-effect pragmas, 1-167
fwrite() function, 3-22

G

GCC compatibility extensions, 1-187
GCC compatibility mode, 1-187
gen_bartlett (generate bartlett window) function, 4-98
gen_blackman (generate blackman window) function, 4-100
general optimization pragmas, 1-159, 1-178
generating
instrumented code, 1-198
gen_gaussian (generate gaussian window) function, 4-101
gen_hamming (generate hamming window) function, 4-102
gen_hanning (generate hanning window) function, 4-103
gen_harris (generate harris window) function, 4-104
gen_kaiser (generate kaiser window) function, 4-105
gen_rectangular (generate rectangular window) function, 4-106
gen_triangle (generate triangle window) function, 4-107
gen_vonhann (generate von hann window) function, 4-109
__getargv function, 1-243
get_default_io_device, 3-42
get_interrupt_info() function, 1-206

getting string containing error message, 3-192
-g (generate debug information) compiler switch, 1-33
global
  adding guard symbol, 1-224
  control variable ___cplb_ctrl, 1-212, 1-217
  guard symbols, 1-217
  variable, 1-273
global-scope constructor, 1-242
globvar global variable, 2-32
gmtime (convert calendar time into broken-down time as UTC) function, 3-113
gmtime function, 3-26, 3-63, 3-151
GNU C compiler, 1-187
graphical character test (See isgraph function)
__GROUPNAME__ macro, 1-58
-guard-vol-loads (guard volatile loads) compiler switch, 1-34

H

handlers
  signal, 1-204
hardware
  disabled loops, 1-236
  flags setting on BF535 processor, 1-117
  loop registers, 1-208
hardware anomaly
  avoiding, 3-7
hardware errors, 1-207
  value, 1-207
hardware loop counters, 1-253
Harris window, 4-104
header
  stop point, 1-175
header file control pragmas, 1-175
header files, 4-4
  C run-time library
    float.h, 3-17
  search for, 1-45
  standard C library, 3-14
header files (new form)
  cassert.h, 3-31
  cctype.h, 3-31
  cerrno.h, 3-31
  cfloat.h, 3-31
  climits.h, 3-31
  clocale.h, 3-31
  cmath.h, 3-31
  csetjmp.h, 3-31
  csignal.h, 3-31
  cstdarg.h, 3-31
  cstddef.h, 3-31
  cstdio.h, 3-31
  cstdlib.h, 3-31
  cstring.h, 3-31
header files (standard)
  assert.h, 3-16
  ctype.h, 3-16
  device.h, 3-16
  device_int.h, 3-16
  errno.h, 3-16
  iso646.h, 3-17
  limits.h, 3-18
  locale.h, 3-18
  math.h, 3-18
  setjmp.h, 3-20
  signal.h, 3-20
  stdarg.h, 3-20
  stddef.h, 3-20
  stdlib.h, 3-24
  string.h, 3-24
  time.h, 3-24

header files (template library)
<algorithm>, 3-31
<deque>, 3-31
<functional>, 3-32
<hash_map>, 3-32
<hash_set>, 3-32
<iterator>, 3-32
<list>, 3-32
<map>, 3-32
<memory>, 3-32
<numeric>, 3-32
<queue>, 3-32
<set>, 3-33
<stack>, 3-33
<utility>, 3-33
<vector>, 3-33
heap, 1-201
base address, 1-248
emptying free-list, 1-250
index, 1-249
memory control, 1-201
section, 1-244
heap_calloc function, 3-114
heap extension routines, 1-245, 1-248
heap_calloc, 1-245
heap_free, 1-245
heap_malloc, 1-245
heap_realloc, 1-245
heap_free function, 3-116
heap functions
calloc, 1-245
free, 1-245
malloc, 1-245
realloc, 1-245
standard, 1-248
heap index, 3-121
heap_init function, 3-118
heap_install function, 3-119
heap interface
alternate, 1-248
standard, 1-248
with multiple heaps, 1-250
heap_lookup function, 3-121
heap_malloc function, 3-123
heap_realloc function, 3-125
heaps
default, 1-245
defining at link time, 1-246
defining at run-time, 1-247
freeing space for, 1-250
lengths of, 1-248
reinitializing, 1-250
heap_space_unused function, 1-249, 3-127
_heap_table table, 1-246
-help (command-line help) compiler switch, 1-34
hexadecimal digit test (See isxdigit function)
hexadecimal floating-point constants, 1-191
-HH (list *.h and compile) compiler switch, 1-34
histogram (histogram) function, 4-110
-H (list *.h) compiler switch, 1-34
__HOSTNAME__ macro, 1-58
HUGE_VAL macro, 3-19
hyperbolic cosine, 3-88
hyperbolic (see cosh, sinh, tanh functions)

I

icache_invalidate routine, 1-215, 3-80
IDDE_ARGS macro, 1-197
idle mode, 1-138
IEEE-754, 1-261
IEEE-754 floating-point format, 1-265
IEEE floating-point support, 1-267
-ieee-fp compiler switch, 1-261
IEEEFP macro, 1-267

-ieee-fp (slow floating point) compiler switch, 1-36
IEEE single/double precision description, 1-261
ifft2d (n x n point 2-d inverse input fft) function, 4-116
ifft (n point radix 2 inverse fft) function, 4-112
ifftrad4 (n point radix 4 inverse input fft) function, 4-114
-I (include search directory) compiler switch, 1-45
iirdf1 (direct form I impulse response filter) function, 4-121
iirdf1_fr16 function, 4-121
iirdf1_init macro, 4-121
iir_fr16 function, 4-118
iir (infinite impulse response filter) function, 4-118
iir_init macro, 4-118
-i (less includes) compiler switch, 1-36
IMASK
  macro, 1-63
  value, 1-138
implicit inclusion, 1-176
-include (include file) compiler switch, 1-36
indexed
  array, 2-18
  initializer support, 1-100
induction variables, 2-28
infinite impulse response (IIR) filter, 4-118
_init_devtab routine, 1-240
init function, 3-36
initialization
  CPLB, 1-240
  device, 1-240
  memory, 1-239
initializer
  aggregate, 1-100
  indexed support, 1-100
inline
  asm statements, 2-20
  avoiding code, 2-41
  keyword, 2-20, 2-41
inline assembly language support keyword (asm), 1-82
  constructs
    optimization, 1-94
    template, 1-84
    template operands, 1-87
  constructs with multiple instructions, 1-93
inline functions, 3-4
inline function support keyword (inline), 1-79, 1-81
inlining
  automatic, 2-19
  file position, 2-63
  function, 2-19
inner loop, 2-28
input operand, 1-85, 1-94
installation location, 1-48
instantiation
  template functions, 1-173
instruction memory, 3-145
instrumented-code profiling, 1-242
instrumented-code profiling library, 1-242
integer
  data type, 1-261
integer data type, 1-261
interfacing C/C++ and assembly (See mixed C/C++/assembly programming)
interpolation filter, 4-94
interprocedural analysis (IPA), 2-7
  compiler switch, 1-36
interprocedural optimizations, 1-74

interrupt
  function, 3-128
  handler pragmas, 1-156
  service routines (ISR), 1-201
INTERRUPT_BITS default interrupt mask, 1-235
interrupt dispatcher
  cycle count, 1-208
interrupt handler
  re-entrant, 1-202
interrupt_info structure, 1-205
interrupt_kind enumeration, 1-204
interrupt_reentrant pragma, 1-156
interrupts
  enabling, 1-240
  general-purpose, 1-202
interrupt-safe functions, 3-26
Interrupt Service Routines (ISR), 1-64, 1-201
invalidate parameter, 1-221
I/O
  extending to new devices, 3-35
  functions, 3-21
  primitives, 3-35, 3-43
  support for new devices, 3-35
iomanip header file, 3-28
iomanip.h header file, 3-33
I/O primitives, 3-34
  data packing, 3-44
  data structure, 3-45
iosfwd header file, 3-28
ios header file, 3-28
iostream header file, 3-29
iostream.h header file, 3-33
-ipa (interprocedural analysis) compiler switch, 1-36, 2-7
isalnum (detect alphanumeric character) function, 3-130
isalpha (detect alphabetic character) function, 3-131
iscntrl (detect control character) function, 3-132
isdigit (detect decimal digit) function, 3-133
isgraph (detect printable character) function, 3-134
isgraph function, 3-134
isinf (test for infinity) function, 3-135
islower C type function, 3-135
islower (detect lowercase character) function, 3-137
isnan (test for NAN) function, 3-138
iso646.h (Boolean operator) header file, 3-17
isprint (detect printable character) function, 3-140
ispunct (detect punctuation character) function, 3-141
isr-imask-check workaround, 1-63, 1-203
ISRs
  default, 1-205
  defining, 1-202
  registering, 1-204
  saved registers, 1-207
ISRs (see interrupt service routines), 1-201
isr-ssync workaround, 1-64
isspace (detect whitespace character) function, 3-142
-I (start include directory) compiler switch, 1-34
-I- (start include directory list) compiler switch, 1-35
istream header file, 3-29
isupper (detect uppercase character) function, 3-143
isxdigit (detect hexadecimal digit) function, 3-144
IVG15 mode, lowest priorty mode, 1-241

J

-jcs2l compiler switch, 1-37
-jcs2l+ compiler switch, 1-37
-jump-constdata|data|code compiler switch, 1-37
jump tables
  storing, 1-37

K

Kaiser window, 4-105
keywords
  extension, 1-79
keywords (compiler)
  extensions, 1-42
  (See also compiler C/C++ extensions)

L

l1_memcpy function, 3-145
L1 SRAM memory, 1-208
L2 SRAM memory
  caching, 1-208
labs (long integer absolute value) function, 3-147
__LANGUAGE_C macro, 1-226
language extensions (compiler) (See compiler C/C++ extensions)
LC_COLLATE locale category, 3-221
ldexp (exponential, multiply) functions, 3-148
LDF
  migrating from previous VisualDSP++ versions, 1-223
  modifications, 1-223
LDF basic configurations, 1-216
LDF output sections, 3-13
ldiv (long division) function, 3-149
ldiv_t type, 3-149
leaf functions, 1-32, 1-42
libetsi532co.dlb library, 1-117
libetsi535co.dlb library, 1-117
libetsi*.dlb library, 1-116
libetsi.h header file, 1-117
libetsi.h system header file, 1-117
___lib_prog_term label, 3-104
libraries
  C/C++ run-time, 3-3
  DSP run-time, 4-2
library
  calling functions, 3-3
  C run-time reference, 3-53 to 3-231
  format for DSP run-time, 4-43
  functions, documented, 3-49
  linking functions, 3-5
  source code
    working with, 4-3
library placement, 3-9
library source file
  devtab.c, 3-41
limits.h header file, 3-18
line breaks
  in string literals, 1-192
__LINE__ macro, 1-227
Linker Description File (LDF), 1-58
linking
  library functions, 3-5
  pragmas for, 1-160
little-endian, 1-136
llabs function, 3-147
llcountones function, 4-84
lldiv function, 3-149
-L (library search directory) compiler switch, 1-37
-l (link library) compiler switch, 1-38, 1-45
locale.h header file, 3-18
localtime (convert calendar time into broken-down time) function, 3-151
localtime function, 3-26, 3-63, 3-113
locking
  function, 3-57

locking routines
ADSP-BF561 processor, A-31
log10 (base 10 logarithm) function, 3-154
log (log base e) functions, 3-153
long
data type, 1-261
division, 3-149
jump (See longjmp, setjmp functions)
latencies, 2-34
long double
representation, 3-209
long file names
handling with -write-files switch, 1-65
long int storage format, 1-261
longjmp (second return from setjmp) function, 3-155
long long
data type, 1-261
loop
control variables, 2-31
cycle count, 2-59
exit test, 2-31
flattening, 2-69
identification annotation, 2-58
iteration count, 2-47
optimization, 1-156, 2-47
parallel processing, 1-159
resource usage, 2-59
rotation by hand., 2-27
short, 2-25
trip count, 2-65
unrolling, 2-25
vectorization, 2-48
vectorizing, 1-156
loop-carried dependency, 2-26, 2-27
loop counters
hardware, 1-253
loop trip count, 2-30
lower case (See islower, tolower functions)
lvalue
GCC generalized, 1-190
generalized, 1-190

M

__MACHINE__ macro, 1-58
macros
multi-statement, 1-228
predefined, 1-225
__RTTI, 1-70
variable argument, 1-192
writing, 1-228
_main
calling, 1-243
_main function, 1-233
malloc (allocate memory) function, 3-157
map files, .XML files, 1-40
-map (generate a memory map) compiler switch, 1-40
mark registers, 1-241
math functions, 4-12
ceil, 3-84
cosh, 3-88
exp, 3-105
fabs, 3-106
floor, 3-107
fmod, 3-110
ldexp, 3-148
library, 4-12
log, 3-153
modf, 3-165
sinh, 3-182
tanh, 3-223
math.h header file, 3-49, 4-12
math.h (mathematics) header file, 3-18
matrix functions, 4-16
matrix.h header file, 4-16
MAXHEAPS macro, 1-246
maximum performance., 2-40
max (maximum) function, 4-125

-MD (make and compile) compiler switch, 1-39
-MDUSE_CACHE flag, 1-216
-MDUSE_SDRAM flag, 1-216
mean (mean) function, 4-126
MEM_ARGV memory section, 1-197
memchr (find first occurrence of character) function, 3-158
memcmp (compare objects) function, 3-159
memcpy (copy characters from one object to another) function, 3-160
memcpy function, 3-145
memcpy_l1 function, 3-145
-mem (invoke memory initializer) compiler switch, 1-40
memmove (copy characters between overlapping objects) function, 3-161
memmove function, 1-57
memory
  allocation functions, 1-244, 3-24
  changing allocation in, 3-172
  controlling size of, 1-201
  data placement in, 2-21
  initialization, 1-239
  initializer support files, 3-6
  map, 1-201
  maximum performance, 2-21
  protection hardware, 1-208
memory bank pragmas, 1-181
memory-mapped registers (MMR), 1-63, 1-65, 1-206
memory sections, 1-243
  bsz, 1-244
  constdata, 1-244
  cplb_code, 1-211, 1-243
  cplb_data, 1-244
  data1, 1-244
  heap, 1-244
  program, 1-243
  stack, 1-244
  sysstack, 1-244
  voldata, 1-244
memory (See calloc, free, malloc, memcmp, memcpy, memset, memmove, memchar, realloc functions)
memset (set range of memory to a character) function, 3-162
_mi_initialize function, 1-239
minimum code size, 2-40
min (minimum) function, 4-127
misaligned_load built-in functions, 1-147
misaligned memory access, 1-154
misaligned_store built-in functions, 1-147
missing operands
  in conditional expressions, 1-190
mixed C/assembly naming conventions, 1-273
mixed C/assembly programming
  arguments and return, 1-257
  asm() constructs, 1-82, 1-84, 1-87, 1-93, 1-94
  data storage and type sizes, 1-261
  scratch registers, 1-252
  stack registers, 1-253
  stack usage, 1-253
mixed C/assembly reference, 1-272
mixed C/C++/assembly programming, 1-230
mixed C/C++/assembly reference, 1-230
mktime (convert broken-down time into a calendar) function, 3-163
-M (make only) compiler switch, 1-38
-MM (make and compile) compiler switch, 1-39
MMR, 1-63, 1-65, 1-206
mode selection switches, 1-13
modf (modulus, float) functions, 3-165
mon.out file, 1-199, 1-200
  post-processing, 1-200

monstartup routine, 1-242
-Mo (processor output file) compiler switch, 1-39
move memory range (See memmove function)
-MQ (output without compile) compiler switch, 1-39
M_STRLEN_PROVIDED bit, 3-47
-Mt preprocessor switch, 1-39
mu_compress (µ-law compression) function, 4-128
mu_expand (µ-law expansion) function, 4-129
multi-core
  applications, 3-9
  environment, 3-60
  linking, A-14
  private data, 3-60
  processor, 3-59
-multicore compiler switch, 1-40, 3-9
multi-issued instructions, 1-64
multiline asm() C program constructs, 1-93
-multiline compiler switch, 1-40
multiple heaps, 1-245
multiple heap support, 1-250
multiple-instruction asm construct, 1-93
multiple pointer types
  declaring, 2-50
multi-statement macros, 1-228
multi-threaded
  applications, 3-9
  environments, 3-7

N

naming conventions
  C and assembly, 1-273
natural logarithm, 3-153
nCompleted field, 3-48
nDesired field, 3-47
new devices
  I/O support, 3-35
  registering, 3-40
new header file, 3-29
new.h header file, 3-33
newline
  in string literals, 1-40, 1-44
new operator
  with multiple heaps, 1-250
next argument in variable list, 3-227
NMI events, 1-202, 1-207
-no-alttok (disable tokens) compiler switch, 1-41
-no-anach (disable C++ anachronisms) C++ mode compiler switch, 1-69
-no-annotate (disable assembly annotations) compiler switch, 1-41
-no-bss compiler switch, 1-41
__NO_BUILTIN macro, 1-227
-no-builtin (no builtin functions) compiler switch, 1-41
__NO_BUILTIN preprocessor macro, 1-41
-no-circbuf (no circular buffer) compiler switch, 1-42
-no-def (disable definitions) compiler switch, 1-42
-no-demangle (disable demangler) C++ mode compiler switch, 1-69
-no-eh (disable exception handling) C++ mode compiler switch, 1-69
NO_ETSI_BUILTINS macro, 1-118
-no-extra-keywords (not quite -ansi) compiler switch, 1-42
-no-force-link (do not force stack frame creation) compiler switch, 1-42
-no-fp-associative compiler switch, 1-43
-no-full-io compiler switch, 1-43
no implicit inclusion, 1-176

-no-implicit-inclusion C++ mode compiler switch, 1-69
NO_INIT qualifier, 1-163
-no-int-to-fract (disable integer to fractional conversion) compiler switch, 1-43
-no-jcs2l compiler switch, 1-44
-no-jcs2l+ compiler switch, 1-44
-no-mem (not invoking memory initializer) compiler switch, 1-44
-no-multiline compiler switch, 1-44
non-constant initializer support (compiler), 1-102
non-IEEE-754 floating point format, 1-261
non-IEEE-754 floating-point format, 1-261
non-unit stride
  avoiding, 2-30
norm (normalization) function, 4-130
-no-rtti (disable run-time type identification) C++ mode compiler switch, 1-69
-no-saturation (no faster operations) compiler switch, 1-44
-no-std-ass (disable standard assertions) compiler switch, 1-44
-no-std-def (disable standard definitions) compiler switch, 1-45
-no-std-inc (disable standard include search) compiler switch, 1-45
-no-std-lib (disable standard library search) compiler switch, 1-45
__NO_STD_LIB macro, 1-45
-no-threads (disable thread-safe build) compiler switch, 1-45
null pointer, 1-247

O

-Oa (automatic function inlining) compiler switch, 1-46
-O (enable optimization) compiler switch, 1-45, 1-47
OFF cache mode, 1-219
-Ofp (frame pointer optimizations) switch, 1-46
-Og (optimize while preserving debugging information) compiler switch, 1-46
-o (output) compiler switch, 1-47
open function, 3-36
open() function, 3-43
operand constraints, 1-88
optimization
  code, 2-40
  controlling, 1-71
  default, 1-72
  enabling, 1-45
  interprocedural analysis, 1-74
  option, 1-36
  reporting progress, 1-51, 1-52
  switches, 1-45, 2-51
  turning on, 1-75
optimization and debugging performance, 1-46
optimizer, 1-71, 1-156, 2-4
optimizing
  for code size, 2-40
  for speed, 2-40
optimizing asm() C program constructs, 1-94
ostream header file, 3-29
outer loop, 2-28
output operand, 1-85
output operands, 1-94
Overflow
  flag for ETSI functions, 1-117
  global variable, 1-117
overlays, 1-253

-Ov num (optimize for speed versus size) compiler switch, 1-47

P

packed data structures, 1-153
padding
  of struct, 2-9
passing
  arguments, 1-257
  arguments to driver, 1-55
  parameters, 1-257
-path-install (installation location) compiler switch, 1-48
-path-output (non-temporary files location) compiler switch, 1-48
-path-temp (temporary files location) compiler switch, 1-48
-path-tool (tool location) compiler switch, 1-48
-pchdir directory (locate precompiled header repository) compiler switch, 1-49
-pch (recompiled header) compiler switch, 1-49
PCHRepository directory, 1-49
-pedantic (ANSI standard warnings) compiler switch, 1-49
-pedantic-errors (ANSI standard errors) compiler switch, 1-49
peeled iterations, 2-66
peeling amount, 2-66
perror (map error number to error message) function, 3-166
-pguide (profile-guided optimization) compiler switch, 1-50
Pipeline Viewer, 2-34
placement support keyword (section), 1-97
pointer
  arithmetic action on, 1-192
  incrementing, 2-18
  resolving aliasing, 2-32
pointer class support keyword (restrict), 1-80, 1-99
polar (construct from polar coordinates) function, 4-131
-P (omit #line) compiler switch, 1-47
post-processing mon.out file from profiler, 1-199
power (see exp, pow, functions)
pow (raise to a power) function, 3-166
-pplist (preprocessor listing) compiler switch, 1-50
#pragma alignment_region, 1-152
#pragma alignment_region_end, 1-152
#pragma align num, 1-150, 1-157, 2-14
#pragma all_aligned, 2-49
#pragma alloc, 1-167, 2-42
#pragma bank_memory_kind, 1-184
#pragma bank_optimal_width, 1-186
#pragma bank_read_cycles, 1-185
#pragma bank_write_cycles, 1-186
#pragma can_instantiate, 1-175
#pragma code_bank, 1-182
#pragma const, 1-168, 2-43
#pragma data_bank, 1-182
#pragma default_section, 1-162
#pragma diag, 1-178
#pragma different_banks, 1-157, 2-50
#pragma do_not_instantiate instance, 1-175
#pragma hdrstop, 1-175
#pragma instantiate, 1-174, 1-275
#pragma interrupt functions, 3-129
#pragma linkage_name, 1-161
#pragma loop_count, 1-157, 2-47
#pragma no_alias, 1-158, 2-50
#pragma no_implicit_inclusion, 1-176

#pragma no_pch, 1-177
#pragma noreturn, 1-168
#pragma no_vectorization, 1-159, 2-47
#pragma once, 1-177
#pragma optimize_as_cmd_line, 1-160, 1-180
#pragma optimize_for_space, 1-160, 1-180
#pragma optimize_for_speed, 1-160, 1-180
#pragma optimize_off, 1-160, 1-179
#pragma optimize_{off|for_speed|for_space}, 2-46
#pragma pack (alignopt), 1-153
#pragma pad (alignopt), 1-155
#pragma pure, 1-168, 2-43
#pragma regs_clobbered, 1-169, 2-44
pragma result_alignment, 2-44
#pragma result_alignment, 1-173
#pragma retain_name, 1-161
pragmas, 1-148
  alignment_region, 1-152
  alignment_region_end, 1-152
  align num, 1-150, 1-157
  bank_memory_kind, 1-184
  bank_optimal_width, 1-186
  bank_read_cycles, 1-185
  bank_write_cycles, 1-186
  can_instantiate, 1-175
  code_bank, 1-182
  const, 1-168
  data alignment, 1-149
  data_bank, 1-182
  default_section, 1-162
  diag, 1-178
  different_banks, 1-157
  do_not_instantiate instance, 1-175
  exception, 1-156
  function side-effect, 1-167
  hdrstop, 1-175
  header file control, 1-175
  instantiate, 1-174
  interrupt, 1-156
  linkage_name, 1-161
  linking, 1-160
  linking control, 1-160
  loop_count(min, max, modulo), 1-157
  loop optimization, 1-156, 2-47
  memory bank, 1-181
  nmi, 1-156
  no_alias, 1-158
  no_implicit_inclusion, 1-176
  no_pch, 1-177
  noreturn, 1-168
  no_vectorization, 1-159
  once, 1-177
  optimize_as_cmd_line, 1-160, 1-180
  optimize_for_space, 1-160, 1-180
  optimize_off, 1-160, 1-179
  pack (alignopt), 1-153
  pad (alignopt), 1-155
  pure, 1-168
  regs_clobbered string, 1-169
  result_alignment, 1-173
  retain_name, 1-161
  section, 1-162
  stack_bank, 1-183
  symbolic_ref, 1-164
  system_header, 1-178
  template instantiation, 1-173
  vector_for, 1-159
  weak_entry, 1-166
#pragma section, 1-162
#pragma stack_bank, 1-183
#pragma symbolic_ref, 1-164
#pragma system_header, 1-178
#pragma vector_for, 1-159, 2-48
#pragma weak_entry, 1-166

predefined macros, 1-225
__ADSPBLACKFIN__, 1-226
__ADSPLPBLACKFIN__, 1-226
__ANALOG_EXTENSIONS__, 1-226
__cplusplus, 1-226
__DATE__, 1-226
__DOUBLES_ARE_FLOATS__, 1-226
__ECC__, 1-226
__EDG__, 1-226
__EDG_VERSION__, 1-226
__EXCEPTIONS, 1-226
__FILE__, 1-226
__LANGUAGE_C, 1-226
__LINE__, 1-227
__NO_BUILTIN, 1-227
__RTTI, 1-227
__SIGNED_CHARS__, 1-227
__STDC__, 1-227
__STDC_VERSION__, 1-227
__TIME__, 1-227
__VERSION__, 1-227
__VERSIONNUM__, 1-227
prelinker, 1-75
preprocessing
a program, 1-225
preprocessor
generated warnings, 1-195
preserved registers, 1-251
PrimIO device, 3-41
_primio.h header file, 3-45
__primIO label, 3-44
primiolib.c source file, 3-42
primitive I/O functions, 3-45
printable characters, 3-134, 3-140
printable character test (See isprint function)
PRINT_CYCLES(STRING,T) macro, 4-34
printf
extending to new devices, 3-35, 3-41
printf() function, 3-41
procedural optimizations, 1-72
processor
clock rate, 4-40
context on supervisor stack, 1-205
initialization, 3-91
priority level, 1-241
processor counts
measuring, 4-33
processor time, 3-85
-proc (target processor) compiler switch, 1-50
profblkfn.exe program, 1-200
profile-guided optimizations, 1-73, 1-197, 2-6, 2-41
profiling
enabling, 1-242
executable outputs, 1-198
using -p switch, 1-198
with instrumented code, 1-197
profiling library
consuming cycles, 1-200
program control functions
calloc, 3-83
malloc, 3-157
realloc, 3-172
-progress-rep-func compiler switch, 1-51
-progress-rep-gen-opt compiler switch, 1-51
-progress-rep-mc-opt compiler switch, 1-52
progress reporting, 1-51, 1-52
Project Options dialog box, 1-9
Project Wizard, 1-196, 1-209
protection violation exception, 1-219

Q

qsort (quicksort) function, 3-167
QUALIFIER keywords, 1-162

R

raise (raise a signal) function, 3-169
rand function, 3-9, 3-26
random number generator (See rand, srand functions), 3-171
rand (random number generator) function, 3-171
-R- (disable source path) compiler switch, 1-52
read function, 3-38
read/write registers, 1-138
realloc (change memory allocation) function, 3-172
__REALNAME__ macro, 1-58
reciprocal square root (rsqrt) function, 4-141
Rectangular window, 4-106
reductions, 2-26
register
  restoring, 1-255
  saving, 1-254
register_handler, 1-204
register_handler() function, 1-204, 1-205, 3-174
registers
  accumulator, 1-132
  accumulators, 1-208
  arithmetic status, 1-207
  call preserved, 1-252
  circular buffer, 1-208
  clobbered, 1-169
  dedicated, 1-251
  for asm()
    constructs, 1-87
  hardware loop, 1-208
  mark, 1-241
  preserved, 1-251
  reserved, 1-53
  saved during ISR prologue, 1-207
  saved in SYSSTACK, 1-207
  scratch, 1-252
  stack, 1-253
  usage (See mixed C/assembly programming)
regs_clobbered pragma, 1-172
regs_clobbered string, 1-170
remarks
  control pragma, 1-178
remove() function, 3-43
RENAME_ETSI_NEGATE macro, 1-118
rename() function, 3-43
-reserve (reserve register) compiler switch, 1-53
reset addres, 1-235
restricted pointers, 2-32
restrict keyword, 1-99, 2-33
restrict qualifier, 2-32
restrict (See pointer class support keyword (restrict))
retain_name pragma, 1-161
return
  values, 1-259
  value transfer, 1-257
returning
  long integer absolute value, 3-147
revisions of silicon, 1-55
rfft2d (n x n point 2-d real input fft) function, 4-138
rfft (n point radix 2 real input fft) function, 4-134
rfftrad4 (n point radix 4 real input fft) function, 4-136
rms (root mean square) function, 4-140
RND_MOD bit, 1-106, 1-125
RND_MOD flag, 1-118
root mean square (rms) function, 4-140
ROT13 algorithm, A-21
rounding
  n ETSI functions, 1-125

-R (search for source files) compiler switch, 1-52
rsqrt (reciprocal square root) function, 4-141
-rtti (enable run-time type identification) C++ mode compiler switch, 1-70
__RTTI macro, 1-70, 1-227
running
  the executable, 1-198
run-time
  disabling type identification, 1-69
  enabling type identification, 1-70
  stack, 1-244, 1-253
run-time environment, 1-230
  programming (See mixed C/C++/assembly programming)
  (See also mixed C/C++/assembly programming)
RUNTIME_INIT qualifier, 1-163

S

-sat32 (32 bit saturation) compiler switch, 1-53
-sat40 (40 bit saturation) compiler switch, 1-53
saturation
  32-bit, 1-53
  40-bit, 1-53
SAVE_REGS() compiler macro, 1-206
SAVE_REGS(interrupt_info *) macro, 1-207
-save-temps (save intermediate files) compiler switch, 1-54
_Sbrk() library function, 1-201
scratch registers, 1-252
SDRAM
  active, 1-54
-sdram compiler switch, 1-54
search
  character string (See strchr, strrchr functions)
  memory, character (See memchar function)
  path for include files, 1-34
  path for library files, 1-37
section
  elimination, 2-40
  placement, 3-9
  qualifiers, 1-162
-section id (data placement) compiler switch, 1-54
  controlling default names, 1-98
section() keyword, 1-80, 1-97
sections
  placing symbols in, 1-162
SECTSTRING double-quoted string, 1-162
seek function, 3-39
segment (See Placement support keyword (section))
sending
  signal to executing program, 3-169
SET_CLOCK_SPEED macro, 1-235
SET_CLOCK_SPEED value, 1-238
set_default_io_device, 3-42
__SET_ETSI_FLAGS macro, 1-119, 1-124
setjmp (define run-time label) function, 3-176
setjmp.h header file, 3-20, 3-49
set jump (See longjmp, setjmp functions)
setting
  compiler functional options, 1-9
  range of memory to a character, 3-162
  register, 1-235
  start, 1-235

short
  data type, 1-261
  storage format, 1-261
shortfract class, 1-128
shortfract header file, 3-29
-show (display command line) compiler switch, 1-55
SIG_DFL function, 3-177
SIG_IGN function, 3-177
signal
  handling, 1-205
  handling functions, 3-20
  processing transformations, 4-8
signal (define signal handling) function, 3-177
signal.h header file, 3-49
signal.h (signal handling) header file, 3-20
signbits, 1-64
signbits instruction, 1-64
-signed-bitfield (make plain bitfields signed) compiler switch, 1-55
-signed-char (make char signed) compiler switch, 1-55
__SIGNED_CHARS__ macro, 1-60, 1-227
  defining, 1-55
silicon revision, 1-55
__SILICON_REVISION__ macro, 1-56
silicon revision setting, 1-55
sind function, 3-179
sine, 3-179
sinf function, 3-179
sin_fr16 function, 3-179
single case range, 1-193
sinh (sine hyperbolic) functions, 3-182
sin (sine) function, 3-179
-si-revision (silicon revision) compiler switch, 1-55
sizeof operator, 1-192
size qualifiers, 3-22
slotID pointer, 3-60, 3-61
source code
  DSP run-time library, 4-3
space allocator, 1-137
space_unused function, 1-248, 3-183
sqrtd function, 3-184
sqrtf function, 3-184
sqrt_fr16 function, 3-184
sqrt (square root) function, 3-184
square root, 3-184
srand function, 3-26
srand (random number seed) function, 3-185
-S (stop after compilation) compiler switch, 1-53
sstream header file, 3-29
-s (strip debug information) compiler switch, 1-53
ssync() function, 1-138
SSYNC instruction
  avoiding silicon anomalies, 1-64
stack, 1-253
  managing, 1-253
  pointer, 1-237, 1-253, 1-254
  pointer dedicated register, 1-251
  registers, 1-253
  user pointer, 1-237
stack frame
  creating, 1-32, 1-42
  linking, 1-254
  unlinking, 1-255
stack frame chain
  terminating, 1-242
Standard C Library, 3-30
standard I/O streams, 3-34

Standard library functions
abort, 3-54
abs, 3-55
acos, 3-56
adi_core_id, 3-59
asin, 3-65
atan, 3-66
atan2, 3-67
atexit, 3-68
atoi, 3-72
atol, 3-73
atoll, 3-77
bsearch, 3-78
calloc, 3-83
cos, 3-86
div, 3-101
exit, 3-104
free, 3-111
frexp, 3-112
heap_calloc, 3-114
heap_free, 3-116, 3-118
heap_install, 3-119
heap_lookup, 3-121
heap_malloc, 3-123
heap_realloc, 3-125
heap_space_unused, 3-127
isalnum, 3-130
isalpha, 3-131
iscntrl, 3-132
isdigit, 3-133
isgraph, 3-134
islower, 3-137
isprint, 3-140
isspace, 3-142
isupper, 3-143
isxdigit, 3-144
labs, 3-147
ldiv, 3-149
log10, 3-154
longjmp, 3-155
malloc, 3-157
memchr, 3-158
memcmp, 3-159
memcpy, 3-160
memmove, 3-161
memset, 3-162
pow, 3-166
qsort, 3-167
raise, 3-169
rand, 3-171
realloc, 3-172
setjmp, 3-176
signal, 3-177
sin, 3-179
space_unused, 3-183
sqrt, 3-184
srand, 3-185
strbrk, 3-201
strcmp, 3-188
strcoll, 3-189
strcpy, 3-190
strcspn, 3-191
strerror, 3-192
strncat, 3-198
strncmp, 3-199
strncpy, 3-200
strrchr, 3-202
strspn, 3-203
strstr, 3-204
strtok, 3-211
strtol, 3-213
strtoll, 3-215
strtoul, 3-217
strtoull, 3-219
strxfrm, 3-221
tan, 3-222
tolower, 3-225
toupper, 3-226
va_arg macro, 3-227
va_end macro, 3-230

va_start macro d, 3-231
standard math functions, 3-5
start code label, 1-235
START_CYCLE_COUNT macro, 4-33
startup code, 1-196, 1-233
ADSP-BF561 processor, A-12
start-up files, 3-8
statement expression, 1-187
statistical
functions, 4-22
profiling, 2-5
stats.h header file, 4-22
status argument, 3-104
stdard.h header file, 3-20
stdarg.h header file, 3-49
__STDC__ macro, 1-227
__STDC_VERSION__ macro, 1-227
stddef.h header file, 3-20
stderrfd function, 3-40
stdexcept header file, 3-29
stdinfd function, 3-40
stdio.h header file, 3-21, 3-34, 3-49
stdlib.h, 3-49
stdlib.h header file, 3-24
stdoutfd function, 3-40
std references, 3-30
-steee-fp compiler switch, 1-267
sti() function, 1-138
STOP_CYCLE_COUNT macro, 4-33
stop (See atexit, exit functions)
storage formats
short, 1-261
strcat (concatenate strings) function, 3-186
strchr (find first occurrence of character in string) function, 3-187
strcmp (compare strings) function, 3-188
strcoll (compare strings) function, 3-189
strcpy (copy from one string to another) function, 3-190
strcspn (compare string span) function, 3-191
streambuf class, 3-34
streambuf header file, 3-30
strerror function, 3-192
strerror (get string containing error message) function, 3-192
strftime (format a broken-down time) function, 3-193
string
converting to double, 3-205
converting to float, 3-207
literals with line breaks, 1-192
transforming with LC_COLLATE, 3-221
string conversion (See atof, atoi, atol, strtok, strtol, strxfrm functions)
string functions
memchar, 3-158
memcmp, 3-159
memcpy, 3-160
memmove, 3-161
memset, 3-162
strcat, 3-186
strchr, 3-187
strcoll, 3-189
strcpy, 3-190
strcspn, 3-191
strerror, 3-192
strlen, 3-197
strncat, 3-198
strncmp, 3-199
strncpy, 3-200
strpbrk, 3-201
strrchr, 3-202, 3-203
strspn, 3-203
strstr, 3-204
strtok, 3-211
strxfrm, 3-221
string header file, 3-30

string.h header file, 3-24, 3-49
string literals
  multiline, 1-40, 1-44
strings
  converting to long double, 3-209
string-to-numeric conversions, 3-24
strlen (string length) function, 3-197
strncat (concatenate characters from one string to another) function, 3-198
strncmp (compare characters in strings) function, 3-199
strncpy (copy characters from one string to another) function, 3-200
strong entry, 1-29
strpbrk (find character match in two strings) function, 3-201
strrchr (find last occurrence of character in string) function, 3-202
strspn (length of segment of characters in both strings) function, 3-203
strstr (compare string, string) function, 3-204
strstream header file, 3-30
strstr (find string within string) function, 3-204
strtod (convert string to double) function, 3-205
strtof (convert string to float) function, 3-207
strtok (convert string to tokens) function, 3-211
strtok function, 3-9, 3-26
strtol (convert string to long integer) function, 3-213
strtold (convert string to long double) function, 3-209
strtoll (convert string to long long integer) function, 3-215
strtoul (convert string to unsigned long integer) function, 3-217
strtoull (convert string to unsigned long long integer) function, 3-219
struct
  assignment, 1-56
  copying, 1-56
  optimizing, 2-9
-structs-do-not-overlap compiler switch, 1-56
struct tm, 3-24
strxfrm (transform string using LC_COLLATE) function, 3-221
switch statements
  in C, 1-37
symbols
  placing in sections, 1-162
synchronization, A-11, A-17, A-31
  functions, 1-138
  lock variables, A-17
-syntax-only (only check syntax) compiler switch, 1-57
SYSCFG, 1-235
SYSCFG_VALUE nitialization value, 1-235
-sysdef (system definitions) compiler switch, 1-57
sysreg.h header file, 2-37
sysreg_read64 function, 1-138
sysreg_read function, 1-138
sysreg_write64 function, 1-138
sysreg_write function, 1-138
system built-in functions
  , 1-137
  idle mode, 1-138
  IMASK, 1-138
  interrupts, 1-138
  read/write registers, 1-138
  stack space allocation, 1-137
  synchronization, 1-138
  system register values, 1-137
System Configuration register, 1-235

__SYSTEM__ macro, 1-58
system registers, 1-138

T

tand function, 3-222
tanf function, 3-222
tan_fr16 function, 3-222
tangent, 3-222
tanh (hyperbolic tangent) functions, 3-223
tan (tangent) function, 3-222
template
  class, 1-275
  function, 1-275
  instantiations, 1-275
  support in C++, 1-275
template for asm() in C programs, 1-84
template inclusion
  control pragma, 1-176
template instantiation pragmas, 1-173
terminate (See atexit, exit functions)
termination functions, 3-24
testset() built-in function, A-31
TESTSET instruction, A-3, A-31
third-party I/O library, 3-21
thread-safe code, 1-58
thread-safe functions, 3-26
-threads (enable thread-safe build) compiler switch, 1-58
threads flag, 1-58
time (calendar time) function, 3-224
time.h header file, 3-24, 4-40, 4-41
time information, 3-24
__TIME__ macro, 1-227
time_t data type, 3-24, 3-224
-time (tell time) compiler switch, 1-59
time zones, 3-24
-T (linker description file) compiler switch, 1-58
tokens, string convert (See strtok function)
tolower (convert from uppercase to lowercase) function, 3-225
toupper (convert characters to upper case) function, 3-226
transferring
  function arguments and return value, 1-257
transformation functions, 4-8
trip count
  loop, 2-65
  minimum, 2-47
true (See Boolean type support keywords (bool, true, false))
twidfft2d function, 4-148
twidfftf_fr16 function, 4-146
twidfftrad2 function, 4-142
twidfftrad4 function, 4-144
type cast, 1-193
typeof keyword, 1-189

U

UNASSIGNED_FILL
  macro, 1-241
  value, 1-241
UNASSIGNED_VAL bit pattern, 1-235
unclobbered registers, 1-171
uninitialized global variable definitions, 1-29
UNIX signal() function, 1-204
-unsigned-bitfield (make plain bitfields unsigned) compiler switch, 1-59
-unsigned-char (make char unsigned) compiler switch, 1-60
untestset() built-in function, A-31
uppercase characters, 3-143
upper case (See isupper, toupper functions)
user
  identifier, 1-245
USER_CRT macro, 1-232
__USERNAME__ macro, 1-58

-U (undefine macro) compiler switch, 1-28, 1-59

V

va_arg (get next argument in variable list) function, 3-227
va_arg macro, 1-258, 3-227
va_end macro, 3-230
va_end (reset variable list pointer) function, 3-230
variable
  statically initialized, 2-13
variable argument function, 1-258
variable argument macros, 1-192
variable-length argument list
  finishing, 3-230
  initializing, 3-231
variable length array, 1-102
variable length arrays, 1-102, 1-191
variadic function, 1-258
variance (VAR) function, 4-150
var (variance) function, 4-150
va_start macro, 3-231
va_start (set variable list pointer) function, 3-231
VDK semaphores, 3-34
vector functions, 4-24
vector.h header file, 4-24
vectorization
  annotations, 2-71
  avoiding, 2-47
  factor, 2-66
  loop, 2-48
  unroll and jam, 2-66
-verbose (display command line) compiler switch, 1-60
-version (display version) compiler switch, 1-60
__VERSION__ macro, 1-227
__VERSIONNUM__ macro, 1-227
video.h header file, 1-140
video operations, 1-140
  accumulator extract with addition, 1-144
  align operations, 1-141
  disaligned loads, 1-142
  dual 16-bit add or clip, 1-143
  packing, 1-141
  quad 8-bit add subtract, 1-142
  quad 8-bit average, 1-143
  subtract absolute accumulate, 1-145
  unpacking, 1-142
virtual function lookup tables, 1-54, 1-98
VisualDSP++
  compiler (ccblkfn), 1-2
  debugger, 1-33
  IDDE, 1-3, 1-9
  simulator, 3-16, 3-23, 3-35, 3-43
Viterbi decoder, 1-132
Viterbi functions, 1-132
  lvitmax1x16(), 1-133
  lvitmax2x16(), 1-133
  rvitmax1x16(), 1-133
  rvitmax2x16(), 1-133
volatile and asm() C program constructs, 1-94
volatile loads
  disables interrupts during, 1-34
von Hann window, 4-109
vtbl section identifier, 1-98
-v (version & verbose) compiler switch, 1-60

W

warning messages, 1-195
warnings
  control pragma, 1-178
-Warn-protos (warn if incomplete prototype) compiler switch, 1-62
wb-dcache workaround, 1-64
-wb_wt_fix switch, 1-64

wchar_t data type, 3-22
-w (disable all warnings) switch, 1-61
weak entry, 1-29
-Werror-limit (maximum compiler errors) compiler switch, 1-61
white space character test (See isspace function)
window
  functions, 4-31
  generators, 4-31
window.h header file, 4-31
_wordsize.h header file, 3-45
__WORKAROUND_ASTAT_RND_MOD macro, 1-65
__WORKAROUND_AVOID_DAG1 macro, 1-62
__WORKAROUND_AVOID_DAG_LOAD_REUSE macro, 1-63
__WORKAROUND_CSYNC macro, 1-63
__WORKAROUND_IMASK_CHECK macro, 1-63
workaround instructions, 1-54
__WORKAROUND_KILLED_MMR_WRITE macro, 1-65
workarounds
  all, 1-62
  astat-rnd_mod, 1-65
  avoid-dag1, 1-62
  avoid-dag-load-reuse, 1-63
  csync, 1-63
  for anomalies, 1-62
  isr-imask-check, 1-63, 1-203
  isr-ssync, 1-64
  killed-mmr-write, 1-65
  sdram-mmr-read, 1-65
  signbits, 1-64
  wb-dcache, 1-64
  wt-dcache, 1-64
__WORKAROUND_SDRAM_MMR_READ macro, 1-65
__WORKAROUNDS_ENABLED macro, 1-62, 1-63, 1-64
__WORKAROUND_SIGNBITS macro, 1-64
__WORKAROUND_SSYNC macro, 1-64
_WORKAROUND_WB_DCACHE macro, 1-64
-workaround (workaround id) compiler switch, 1-62
  all, 1-62
  astat-rnd_mod option, 1-65
  avoid-dag1 optiion, 1-62
  avoid-dag-load-reuse option, 1-63
  csync option, 1-63
  isr-imask-check option, 1-63
  isr-ssync option, 1-64
  killed-mmr-write option, 1-65
  sdram-mmr-read option, 1-65
  signbits option, 1-64
  wb-dcache option, 1-64
  wt-dcache option, 1-64
__WORKAROUND_WT_DCACHE macro, 1-64
-W (override error) compiler switch, 1-60
-Wremarks (enable diagnostic warnings) compiler switch, 1-61
Write Back mode, 1-218
-write-files (enable driver I/O pipe) compiler switch, 1-65
write function, 3-37, 3-38
-write-opts (enable driver I/O pipe) compiler switch, 1-66
writes
  array element, 2-28
Write Through mode, 1-214, 1-218
writing
  preprocessor macros, 1-228

-Wterse (enable terse warnings) compiler switch, 1-61

X

.XML files, 1-40
-xref (cross-reference list) compiler switch, 1-66

Z

zero_cross (count zero crossing) function, 4-151
ZERO_INIT qualifier, 1-163
zero length arrays, 1-191